Numerical Methods in Global Transport-Chemistry Models

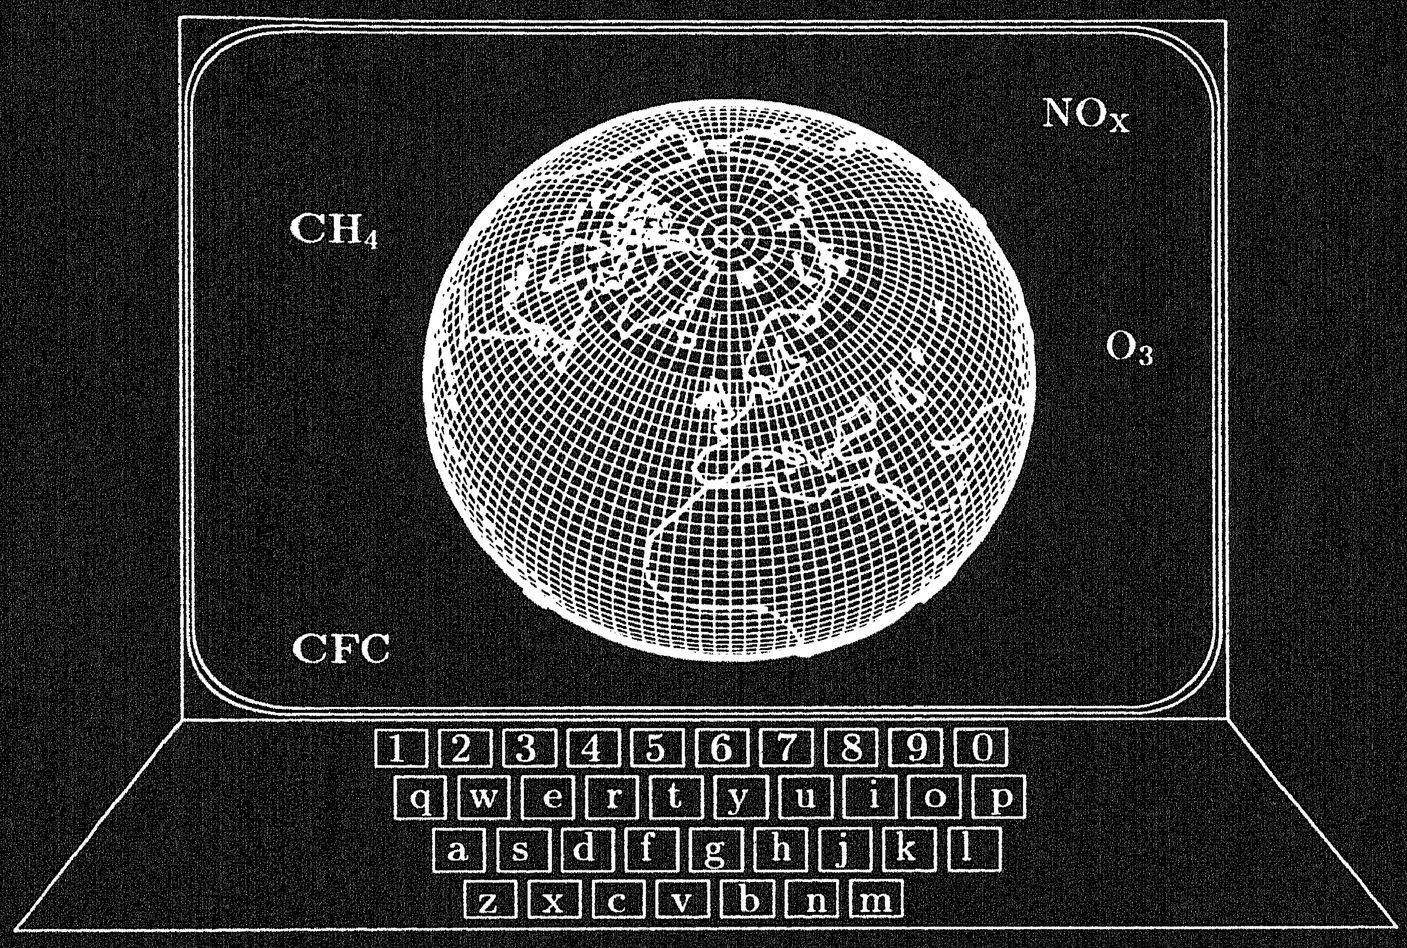

Edwin Spee

Numerical Methods
in
Global Transport-Chemistry Models

ISBN: 90-74795-88-9
Printing & binding: Ponsen en Looijen bv, Wageningen
No. of copies: 300
Typesetting: LaTeX 2.09

Numerical Methods in Global Transport-Chemistry Models

ACADEMISCH PROEFSCHIFT

ter verkrijging van de graad van doctor
aan de Universiteit van Amsterdam,
op gezag van de Rector Magnificus
prof. dr J.J.M. Franse
ten overstaan van een door
het college voor promoties ingestelde commissie
in het openbaar te verdedigen in de Aula der Universiteit
op vrijdag 23 januari 1998 te 15:00 uur

door

Edwin Johannes Spee

geboren te Gramsbergen

Faculteit Wiskunde, Informatica, Natuurkunde en Sterrenkunde

Promotor: prof. dr P.J. van der Houwen
Co-promotoren: dr J.G. Verwer
dr W. Hundsdorfer

Het onderzoek dat tot dit proefschrift heeft geleid werd mede mogelijk gemaakt door financiële steun van het Rijksinstituut voor Volksgezondheid en Milieuhygiene (RIVM) te Bilthoven in het kader van het CIRK-project.

Een OIO zat lang te debuggen
Maar plotseling leek het te lukken
Hij schreef zijn verslag
En toen kwam de dag
Waarop hij zijn proefschrift liet drukken

Voorwoord

De inhoud van dit proefschrift bestaat uit een inleiding en een korte modelbeschrijving gevolgd door vijf min of meer op zich zelf staande artikelen en zijn geschreven samen met diverse collega's.
Het betreft, op enkele kleine wijzigingen na, de volgende artikelen:

- An efficient horizontal advection scheme for modeling of global transport of constituents - W. Hundsdorfer and E.J. Spee. *Mon. Wea. Rev.*, 123:3554–3564, 1995.

- Coupling advection and chemical kinetics in a global atmospheric test model - E.J. Spee. In H. Power, N. Moussiopoulos, and C.A. Brebbia, editors, *Air Pollution III, Volume 1: Air Pollution, Theory and Simulation*, pages 319–326. Computational Mechanics Publications, Southampton, Boston, 1995.

- A numerical study for global atmospheric transport-chemistry problems - E.J. Spee, J.G. Verwer, P.M. de Zeeuw, J.G. Blom, and W. Hundsdorfer. 1997. Revision of MAS-R9702.

- Sensitivity of atmospheric transport model performance to numerical advection schemes and resolution
  - E.J. Spee, A.C. Petersen, H. van Dop, and W. Hundsdorfer. Revision of Report MAS-R9710, CWI, Amsterdam, 1997. Submitted to *J. Geophys. Res.*

- A second-order Rosenbrock method applied to photochemical dispersion problems - J.G. Verwer, E.J. Spee, J.G. Blom, and W. Hundsdorfer. Report MAS-R9717, CWI, Amsterdam, 1997. Submitted to *SIAM J. Sci. Comput.*

Dit proefschrift zou nooit tot stand zijn gekomen zonder de steun van en samenwerking met enkele collega's en begeleiders. Met het gevaar dat ik iemand over het hoofd zie, wil ik toch een aantal namen noemen.

Jan Verwer en Willem Hundsdorfer hebben mij bijgeschoold tot een numeriek wiskundige. Toen ik op het CWI kwam had ik zelfs nog nooit gehoord van stijve differentiaalvergelijkingen en flux-limiters, dus ik had ook nog wel wat te leren. Verder leerden zij mij hoe wetenschappelijke resultaten preciezer te formuleren. Hun deur stond altijd voor mij open.

Piet van der Houwen dank ik voor het feit dat hij dit proefschrift wilde 'promoten'.

Alle co-auteurs van de hierboven genoemde artikelen wil ik bedanken voor de goede samenwerking in de afgelopen vier jaar.

Joke Blom, Jacques de Swart en Walter Stortelder wil ik bedanken voor de diverse UNIX, LaTeXen Fortran tips.

Paul Richardson wil ik bedanken voor het controleren van de hoofdstukken 1, 2 en 4 op Engels taalgebruik.

In de afgelopen vier jaar heb ik drie kamergenoten 'versleten'. Ik wil jullie bedanken voor de goede sfeer, die overigens in heel NW en later MAS heerst. Van Maarten van Loon en Paul de Zeeuw heb ik de kunst van het promoveren afgekeken. Debby Lanser wil ik bedanken voor de kritische blik op enkele onderdelen van dit werk.

Als laatste wil ik mijn ouders, Ingrid, Wilfred en de neefjes Jochem en Olaf, Martine en Peter en mijn vrienden bedanken voor hun morele steun en gezelligheid.

Bergen nh, november 1997
Edwin Spee

Contents

Contents

Chapter 1

Introduction

1.1 Air quality models

This Ph. D. thesis deals with numerical algorithms for use in air quality models, particularly global models. Air quality models are used to enhance the understanding of the chemical composition of the atmosphere, in particular with regard to the cause and effect of air pollution. With global models research is carried out into global environmental problems, such as the ozone hole at the South Pole and the anthropogenically enhanced greenhouse effect resulting in global warming. The ozone hole is caused by emissions of chlorofluorocarbons (CFC) resulting in a strong breakdown in the Antarctic spring of ozone (O_3), and the greenhouse effect by increasing concentrations of methane (CH_4) and carbon dioxide (CO_2) and possibly other trace gases, including for example O_3. A related field of research is smog prediction, a more local phenomenon, where regional air quality models are used.

Examples of global air quality models are Moguntia [Zim87], IMAGES [MB95] and TM3. TM3 is currently under development in the Dutch Centre for Climate Research (CKO), and is the successor of TM and TM2 [HK89]. Examples of regional models are CWIROS [vL96] and LOTOS [BR97].

Air quality models describe movement by wind (advection) of species, diffusion, cumulus cloud convection, chemical reactions, emissions and depositions, see Figure 1.1. These models can be either on-line or off-line. In the on-line approach the transported species are treated as additional variables of a complete atmospheric circulation model and parametrizations in the circulation model can now be affected by the polluted species. In the off-line approach the transported species are treated independently of any other model, and are advected by a wind field stored in a meteorological database. In this thesis we studied numerical algorithms in off-line models, but these techniques can of course also be used in on-line models.

Advection is an important process in air quality models. For example, the CFC

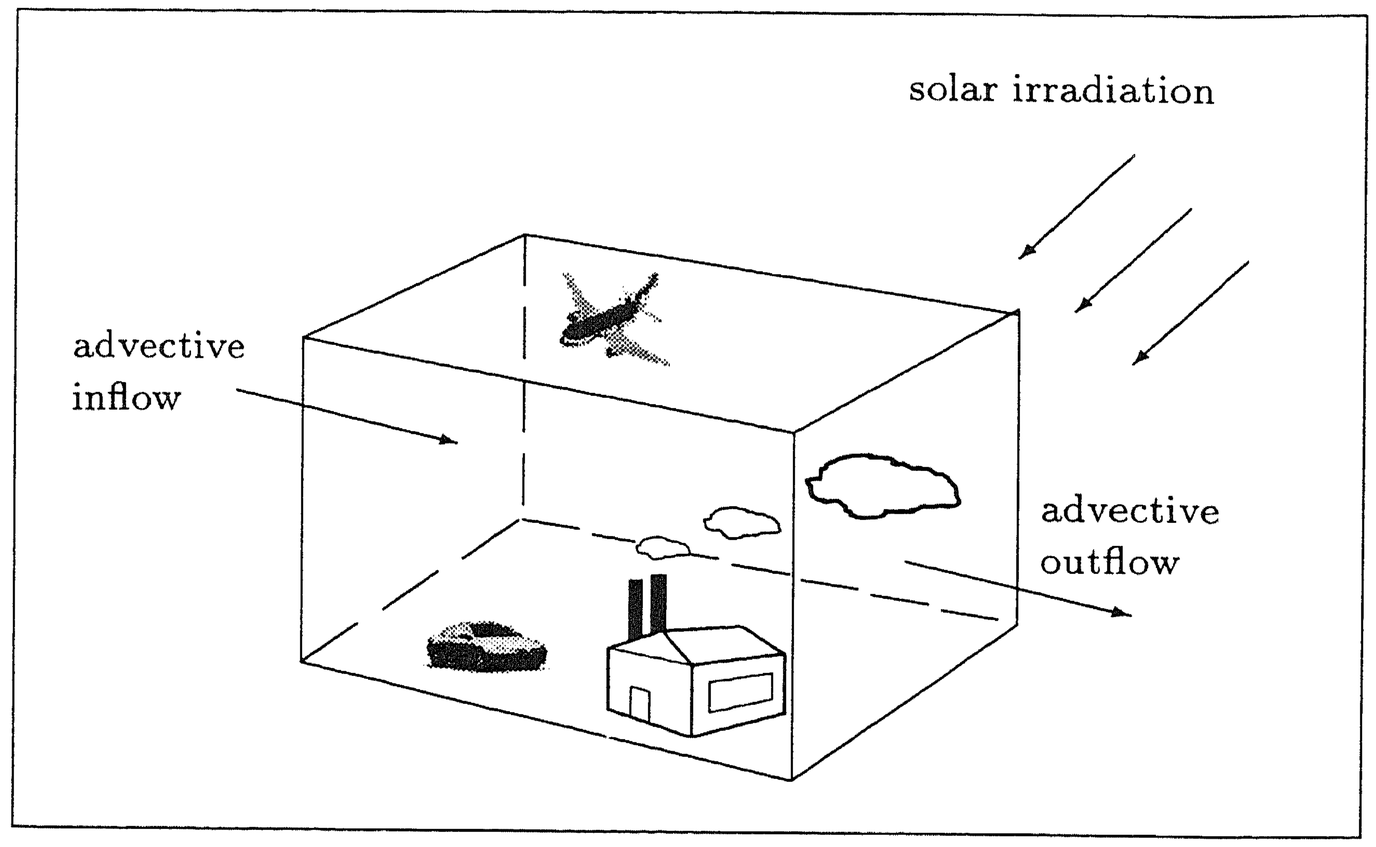

Figure 1.1: An example of a box-model: emission by cars, factories and airplanes and, due to wind and diffusion, exchange of polluted species between neighboring grid cells.

causing the ozone hole at the South Pole are emitted predominantly in the industrialized Northern Hemisphere, and due to advection the CFC does reach the South Pole. On the other hand, wind carries clean sea-air to land, reducing the pollution over land areas.

In the vertical direction, transport is dominated by turbulent diffusion and cumulus cloud convection. Advection is less important due to the low vertical wind components. Without a good description of the vertical transport in the model, emitted species stay at the ground level where they are emitted, and the ground level concentrations will be too high. It also affects horizontal transport, because the horizontal wind velocities are relatively low at ground level, and reach a maximum at medium to high levels depending on meteorological conditions.

A well-known example of chemical reactions is the already mentioned degradation of ozone by CFC at the South Pole. Another example is the formation of ozone, a species which is not emitted itself. In urban regions ozone is formed by a reaction involving nitrogen-oxides (NO_X) and oxygen (O_2). Nitrogen-oxides are emitted by cars and factories, and oxygen is with a mixing ratio of about twenty percent one of the main constituents of atmospheric air. A significant part of the chemical reactions strongly depend on sunlight, causing large differences between day and night.

1.1.1 The CIRK project

The research described in this Ph. D. thesis has been carried out within the CIRK project, a joint research effort of CWI (Center for Mathematics and Computer Science), IMAU (Institute for Marine and Atmospheric Research), RIVM (National Institute of Public Health and Environmental Protection) and KNMI (Royal Netherlands Meteorological Institute). The goal of the project has been the improvement of global 3D modeling of the troposphere including chemistry, source and sink terms, and troposphere/stratosphere exchange processes. The emphasis in this thesis is on numerical and computational aspects of the modeling. In parallel with this research, vertical transport and sub-grid scale processes are studied at IMAU. For the research in this thesis the contact with researchers from IMAU, RIVM and KNMI has been very helpful to focus the effort on problems that arise in the present global models.

1.2 Numerical techniques for air quality models

Modern air quality models are based on mass balances and are solved on an appropriate, Eulerian grid. Standard numerical techniques will not work due to the large number of unknowns and the different time scales in the model. The number of unknowns is the product of the chemical species involved, with the total number of grid points. Current models use approximately 10 to 100 chemical species. Horizontal resolutions vary between $10^\circ \times 10^\circ$ and $0.5^\circ \times 0.5^\circ$ resulting in respectively $36 \times 18 = 648$ and $720 \times 360 = 259200$ horizontal grid points. For a good description in the vertical direction 15 to 60 layers are needed. Hence the number of unknowns approximately varies between 10^5 and 10^9. The time scale varies from less than a second for chemical species involving radicals, to years. These wide ranging time scales lead to stiff chemistry, requiring implicit solvers. It is clear that we need the fastest computers available, such as vector computers, parallel computers or clusters of workstations, combined with fast numerical techniques. To obtain high performance on a vector computer, and good load balancing on parallel computers, it is desirable that a numerical technique can be used with a large, fixed time step equally distributed over the grid. In this application, a large time step means a time step of about twenty minutes.

In [PBC+95] it is pointed out that the success of the next generation of air quality models, in terms of both its scientific and societal impacts, is dependent on how well new computing platforms and methodologies are exploited by the atmospheric science community. A thorough review on the numerical aspects of air quality models can be found in [Wil92, MGS82].

We have to balance discretization errors (in time and space) with errors originating from model errors, such as large uncertainties in, e.g., the chemical rate constants, and the emissions. Another source of model errors is the limited number of chemical species and reactions. Therefore 1% error in the numerics is good enough. However,

budget studies and long term calculations require exact mass conservation.

Most air quality models solve their systems of equations by subdividing the problem in a small number of subprocesses, which now can be solved with different numerical techniques and different time steps. This happens for example with the known operator splitting technique. The model is usually subdivided in three parts, firstly advection, secondly chemistry including emissions and depositions and finally diffusion and sub-grid scale parametrizations such as cumulus cloud convection. We focused on numerical schemes for advection and chemistry, and developed alternatives for operator splitting.

1.2.1 Advection schemes

In the development of advection solvers we focused on mass conservative schemes, which do not generate undershoot or overshoot. Avoiding undershoot is very important, because undershoot may lead to unphysical, negative concentrations, which can give an instability in the chemistry step. Mass conservation is accomplished by formulating advection in terms of in- and outgoing fluxes. The most simple scheme of this type is the so-called Donor Cell scheme, where advection terms are discretized with first-order upwind differences. Unfortunately, this approach leads to a significant amount of numerical (artificial) diffusion, especially on coarse grids. In order to avoid excessive numerical diffusion we use higher order schemes, together with flux-limiting to minimize undershoot and overshoot.

With the used coordinate system we have a singularity at both poles on a sphere. To overcome this difficulty, we focused on two options: a so called 2D spherical reduced grid where less cells are used near the poles, and a scheme which is unconditionally stable in the longitudinal direction.

Alternative advection schemes are sub-grid schemes and semi-Lagrangian schemes. Sub-grid schemes, such as Slopes [RL81] and Second Moment [Pra86], are too expensive in terms of memory. Moreover, using the sub-grid information in the chemistry calculation would lead to a large increase in cpu time, see Chapter 6. Semi-Lagrangian schemes can use very large time steps, but such schemes are not mass conservative. Moreover, in actual models, the other subroutines (chemistry, vertical transport) and their combination via time splitting will prohibit very large time steps anyhow.

1.2.2 Chemistry solvers

Since chemistry often consumes the largest part of the cpu time, efficiency in terms of cpu time is the most important property for chemistry solvers. This can be accomplished by using large time steps, but this in only possible if the method used has good stability properties. Many existing models use dedicated explicit solvers, such as QSSA [HHI78] and EBI [HBCH93]. Black-box stiff ODE solvers are considered too costly. We developed dedicated linearly implicit solvers of Rosenbrock type.

These solvers exploit sparsity in the Jacobian matrix. The newly developed chemistry solvers only need as input the chemical equations with their rate constants, so a change in the chemical system is very quickly processed in the model.

1.2.3 Alternative splitting techniques

The model equation is usually solved with operator splitting. When using operator splitting, different numerical techniques with different, local, time steps can be used for different parts of the model. Furthermore, the building of the computer code of the model is modular and memory efficient.

A disadvantage is the creation of a splitting error. Therefore we developed some alternatives for the commonly used operator splitting. In Chapter 4 we extended Strang splitting [Str68] with the integration of the chemistry along characteristics. This approach is effective only if the reaction constants are updated in each chemistry time step, because the space dependency originates from the reaction constants, which are usually fixed for fairly long times. In Chapter 5 we used Twostep as time integrator for a complete chemistry-diffusion step. In Chapter 7 we give results for some experiments with ROS2 applied with approximate factorization as time integrator for a complete chemistry-diffusion-advection step. The results in both chapters show that diffusion can be coupled with chemistry in an efficient manner. On the other hand, our experiences also indicate that in global models the splitting error is acceptable. More research is needed to come to more definite conclusions on the splitting error.

1.3 Outline of this Ph. D. thesis

The remainder of this Ph. D. thesis consists of a brief model description and five articles. The articles are self-contained.

In Chapter 2 we give a mathematical description of the model and some technical problems arising from coordinate transformation. We finish this chapter with experiences with a state of-the-art chemical kinetics preprocessor.

In Chapter 3 [HS95] a new technique for solving the 2D advection equation on a sphere is described. We combined unconditional stability with mass conservation. Note that we show in Section 6.7 how to use the new scheme in a realistic 3D setting.

In Chapter 4 [Spe95, HS96] a 2D global model including advection and chemistry is developed, and serves as a benchmark for numerical solvers. Chemistry is integrated along characteristics to obtain a highly accurate reference solution for our rather complicated model problem. This approach is also used in Chapter 5. Strang splitting on a reduced grid gives good results.

In Chapter 5 [SVdZ+97a] we present a benchmark problem including horizontal advection, vertical diffusion and chemistry. With this benchmark we compared

two different numerical techniques, with respect to accuracy and performance on a Cray C90.

In Chapter 6 [SPvDH97] we tested several 3D advection schemes. Firstly, we extended the scheme from Chapter 3 to 3D on a reduced grid and named the scheme Split-rg. Other schemes in this comparison are the newly developed Mol-rg, a 3D Method of Lines (MOL) scheme on a reduced grid and the existing schemes Slopes, Second Moment and Donor Cell. We measured accuracy versus efficiency in terms of cpu time using the well-known solid-body rotation on a sphere. Furthermore, the newly developed schemes together with Slopes were implemented in TM2. Using wind fields from a meteorological model we compared radon concentrations from the model with measurements.

In Chapter 7 [VSBH97] we studied ROS2, a second-order Rosenbrock method for solving the chemical rate equations. ROS2 integrates box-models with a fixed time step up to twenty minutes, even in a 3D model with Strang splitting where each advection step introduces artificial transients.

1.4 Summary

In this Ph. D. thesis we present new numerical methods for global air quality models. The performance of these techniques is extensively illustrated through numerical experiments. The experiments where chemistry and advection or diffusion are involved, are newly developed. These new benchmarks are based on commonly used chemical box-models and the solid-body rotation on a sphere. With these experiments we are able to measure both accuracy and efficiency in terms of cpu time for the complete chemistry-advection-diffusion solver.

Existing models use operator splitting for time stepping. Although we observed a small splitting error in Chapter 4, 5 and 7, the second-order Strang splitting seems as the best approach yet in air quality models. Compared to alternatives, its main advantage lies in its flexibility in coding and general applicability.

Accurate results at low computational costs on high performance computers are obtained using a spherical reduced grid for spatial discretization. On this grid, accurate results are obtained with the advection solver Split-rg from Chapter 6. Split-rg is formulated in terms of in- and outgoing fluxes, and therefore mass conservative. Split-rg uses flux-limiting to minimize over- and undershoot.

To obtain high performance in the chemistry part of the model, it is necessary to use a time step which is constant over a large part of the computational domain. This can be accomplished by using clustering as is done in Chapter 5 using Twostep, but it is easier to have the same time step for the whole computational domain, as is done in Chapter 7 using ROS2, a second-order Rosenbrock scheme. The results in Chapter 7 show that ROS2 is able to integrate complex chemical systems with a fixed time step up to twenty minutes, a very good result for a solver which does not use any ad-hoc rule.

Chapter 2

Model description

2.1 Introduction

This chapter gives a basic mathematical description of currently used global transport-chemistry models, such as TM3. These global models use certain simplifications with respect to orography for example, and these simplifications are motivated in this chapter. We start in Section 2.2 with the transport-chemistry equation. In Section 2.3 we give a short introduction to pressure based hybrid η-coordinates, and discuss the transformation of the model equation in this system. Section 2.4 gives some information on the preprocessor KPP used in the chemistry part of our air quality test models.

2.2 Transport-chemistry equation

The mathematical formulation of air quality models leads to the transport-chemistry equation, which describes advection by wind of chemical species, the effects of (turbulent) diffusion, (sub-grid scale) parametrizations, chemical reactions, sources and sinks,

$$\frac{\partial c_l}{\partial t} + \vec{\nabla} \cdot (\vec{u} c_l) = \vec{\nabla} \cdot \left(\rho D \vec{\nabla} \left(\frac{c_l}{\rho} \right) \right) + P_l(c) + F_l(c), \quad l = 1, \ldots, N_c.$$

The concentration c [particles m^{-3}] is a vector with dimension N_c containing all chemical species, $\vec{u}$ is the given wind field [m s^{-1}], ρ is the density [kg m^{-3}], D is a tensor describing diffusion [m^2 s^{-1}], P and F [s^{-1}] describe respectively parametrizations such as cumulus cloud convection, and the chemical reactions, sources and sinks.

In a numerical model, it is important to treat tracer transport in a similar way as the background density, see [Wil92]. Tracer transport should therefore not be

considered independently from the continuity equation for the atmosphere, which reads

$$\frac{\partial \rho}{\partial t} + \vec{\nabla} \cdot (\vec{u}\rho) = 0. \tag{2.1}$$

Here it is assumed that the source and sink terms cancel out in the contribution to the total density. For example, the emission of NO_X is compensated by a decrement of N_2 and O_2.

Using the relation $c = \rho\,\chi$ with χ the mixing ratio [particles / kg air], we rewrite the transport-chemistry equation into

$$\frac{\partial\,(\rho\,\chi_l)}{\partial t} + \vec{\nabla} \cdot (\vec{u}\rho\,\chi_l) = \vec{\nabla} \cdot \left(\rho\,D\,\vec{\nabla}\chi_l\right) + P_l(\rho\,\chi) + F_l(\rho\,\chi), \quad l = 1, \ldots, N_c. \tag{2.2}$$

In Section 6.7 it is shown that this formulation with mixing ratios gives better results in terms of mass conservation and monotony than the formulation with concentrations when using dimension splitting in flux form in the advection part of the model.

2.3 Hybrid η-coordinates

Most recent atmospheric and meteorological calculations are performed with terrain-following coordinates in the vertical direction, and longitude-latitude coordinates (λ, ϕ) on the sphere. The first terrain-following coordinates are σ-coordinates introduced in [Phi57]. A disadvantage of the σ-coordinates is the fact that the effects of the orography exists in all vertical layers up to the highest one. In a general circulation model it is more efficient to have layers above some level with no effect on the orography, because the mathematical description then simplifies in a part of the model. Therefore the European Centre for Medium Range Weather Forecasts at Reading, UK (*ECMWF*) developed a hybrid coordinate system [Lou86], which is terrain-following at the surface of the earth, and has pressure levels at the top of the model. This pressure based coordinate system is defined (strictly speaking: a continuous analogue of the discrete definitions is given) by the relation

$$p(\lambda, \phi, \eta, t) = A(\eta) + B(\eta)\,p_{sfc}(\lambda, \phi, t), \tag{2.3}$$

where η is the vertical coordinate, p is the pressure (Pa), p_{sfc} is the pressure at ground level, t is time. The functions $A(\eta)$ and $B(\eta)$ are given. There are two special cases. If $A = 0$ we have pure σ-coordinates. With $B = 0$ we have pure pressure levels. See Table 2.1 for the discrete 15 and 31 layer models.

	hybrid coordinates				pressure levels			
k	$A_{k+\frac{1}{2}}$	$B_{k+\frac{1}{2}}$	p	z	$A_{k+\frac{1}{2}}$	$B_{k+\frac{1}{2}}$	p	z
	Pa		hPa	km	Pa		hPa	km
0	0.000000	0.0000000000	0	30.94	0	0	0	38.2
1	2000.000000	0.0000000000	20	26.46	1000	0	10	31.2
2	4000.000000	0.0000000000	40	21.99	3000	0	30	24.1
3	6000.000000	0.0000000000	60	19.41	5000	0	50	20.9
4	8000.000000	0.0000000000	80	17.59	7000	0	70	18.7
5	9976.135361	0.0003908582	100	16.17	10000	0	100	16.5
6	11820.539617	0.0029197006	121	14.96	15000	0	150	13.9
7	13431.393926	0.0091941320	144	13.88	20000	0	200	12.1
8	14736.356909	0.0203191555	168	12.89	25000	0	250	10.6
9	15689.207458	0.0369748598	194	11.96	30000	0	300	9.4
10	16266.610500	0.0594876397	223	11.09	40000	0	400	7.4
11	16465.005734	0.0878949492	254	10.25	50000	0	500	5.7
12	16297.619332	0.1220035886	287	9.46	70000	0	700	3.0
13	15791.598604	0.1614415235	321	8.69	85000	0	850	1.3
14	14985.269630	0.2057032385	358	7.95	92500	0	925	0.7
15	13925.517858	0.2541886223	397	7.23	100000	0	1000	0.0
16	12665.291662	0.3062353873	437	6.55				
17	11261.228878	0.3611450218	478	5.89				
18	9771.406290	0.4182022749	521	5.26				
19	8253.212096	0.4766881754	565	4.65				
20	6761.341326	0.5358865832	610	4.07				
21	5345.914240	0.5950842740	656	3.51				
22	4050.717678	0.6535645569	703	2.98				
23	2911.569385	0.7105944258	749	2.47				
24	1954.805296	0.7654052430	795	2.00				
25	1195.889791	0.8171669567	840	1.55				
26	638.148911	0.8649558510	883	1.15				
27	271.626545	0.9077158297	922	0.78				
28	72.063577	0.9442132326	957	0.48				
29	0.000000	0.9729851852	986	0.23				
30	0.000000	0.9922814815	1005	0.07				
31	0.000000	1.0000000000	1013	0.00				

Table 2.1: Constants A and B defining the hybrid coordinates for two different grids: a 15 and a 31 layer model. The values of p and z are for a standard atmosphere. The altitude z is equal to $r - r_{sfc}$.

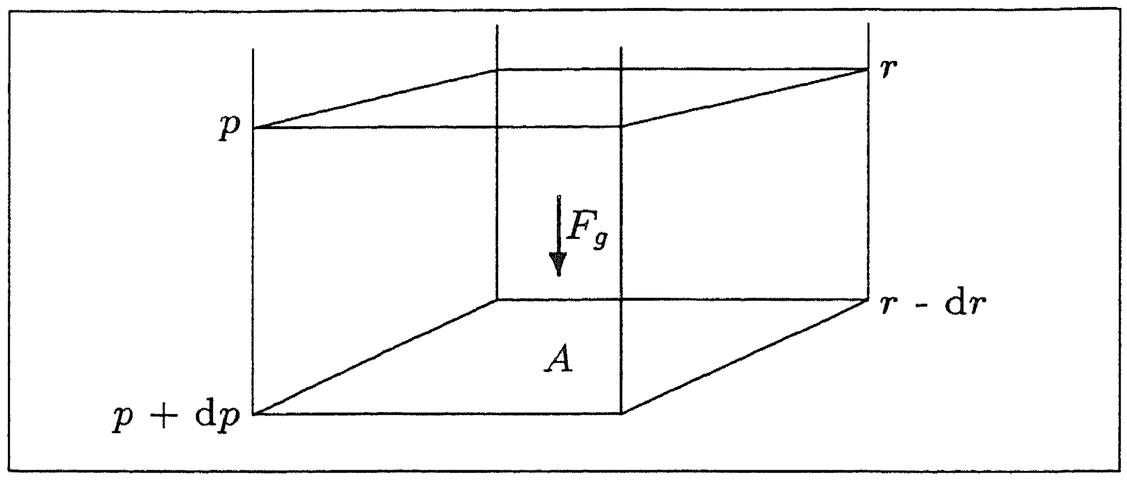

Figure 2.1: Pressure increases by the weight of gas

2.3.1 Coordinate transformation

The coordinate transformation from Cartesian coordinates to hybrid coordinates is given by

$$\begin{aligned} x &= R(\lambda,\phi,\eta,t)\ \cos\phi\ \cos\lambda \\ y &= R(\lambda,\phi,\eta,t)\ \cos\phi\ \sin\lambda \\ z &= R(\lambda,\phi,\eta,t)\ \sin\phi \end{aligned} \tag{2.4}$$

where $\lambda \in [0, 2\pi)$, $\phi \in [-\pi/2, +\pi/2]$ are the longitude and latitude coordinates and $R(\lambda,\phi,\eta,t)$ is a function with boundary conditions $R(\lambda,\phi,0,t) = r_{top}$, the top of the model, and $R(\lambda,\phi,1,t) = r_{sfc}(\lambda,\phi)$, ground level for all λ, ϕ and t. We observe that R depends through the pressure p on λ, ϕ, η and t,

$$R(\lambda,\phi,\eta,t) \;=\; r(p(\lambda,\phi,\eta,t)), \tag{2.5}$$

where p is given by Equation (2.3). The dependence of the radial distance r on p follows from the hydrostatic approximation. As shown in Figure 2.1, the pressure increases with the weight of the gas F_g in the cube with an infinitesimal volume $\mathrm{d}V$, divided by the area A of the cube,

$$\mathrm{d}p = \frac{\mathrm{d}F_g}{A} = \frac{g\,\rho\,\mathrm{d}V}{A} = -g\,\rho\,\mathrm{d}r = -\frac{g\,M_{air}}{R_{gas}}\,\frac{p}{T}\,\mathrm{d}r = -\gamma^{-1}\,\frac{p}{T}\,\mathrm{d}r,$$

with $R_{gas} = 8.3143$ J mol^{-1} K^{-1}, the molar gas constant, $g = 9.81$ m s^{-2}, the acceleration of gravity, $M_{air} = 2.897 \times 10^{-2}$ kg mol^{-1}, the average molecular weight of air [War88] and $\gamma = R_{gas}/(g\,M_{air}) = 2.9256 \times 10^{-2}$ km K^{-1} for dry air. For a constant temperature T we have a closed form for the relation between r and p,

$$r(p) \;=\; r_{sfc} + \int_{p_{sfc}}^{p} -\gamma\,\frac{T}{\Phi}\,\mathrm{d}\Phi \;=\; r_{sfc} + \gamma\,T\,\ln\left(\frac{p_{sfc}}{p}\right). \tag{2.6}$$

Differentiation of Equations (2.3) and (2.6) leads to an expression for the scale factor in the vertical direction h_η,

$$h_\eta = \left|\frac{\partial R}{\partial \eta}\right| = \left|\frac{\partial R}{\partial p}\frac{\partial p}{\partial \eta}\right| = \gamma \frac{T}{p}\left(\frac{\partial A(\eta)}{\partial \eta} + p_{sfc}(\lambda, \phi, t)\frac{\partial B(\eta)}{\partial \eta}\right). \tag{2.7}$$

If we have pure pressure levels and $A(\eta) = \eta$, this leads to

$$h_\eta = \gamma \frac{T}{p}. \tag{2.8}$$

2.3.2 Continuity equation in hybrid η-coordinates

If R does not depend on λ and ϕ, which means that the earth is a perfect sphere, p_{sfc} and T do not depend on λ and ϕ, the continuity equation for the atmosphere (2.1) and model equation (2.2) transforms in hybrid η-coordinates into respectively,

$$\frac{\partial \rho}{\partial t} + \frac{1}{r\,\cos\phi}\left[\frac{\partial(u\,\rho)}{\partial \lambda} + \frac{\partial(v\,\cos\phi\,\rho)}{\partial \phi}\right] + \frac{1}{h_\eta\, r^2}\frac{\partial\left(w\, r^2\,\rho\right)}{\partial \eta} = 0 \tag{2.9}$$

and

$$\begin{aligned}\frac{\partial(\rho\chi)}{\partial t} + \frac{1}{r\,\cos\phi}\left[\frac{\partial(u\,\rho\chi)}{\partial \lambda} + \frac{\partial(v\,\cos\phi\,\rho\chi)}{\partial \phi}\right] + \frac{1}{h_\eta\, r^2}\frac{\partial\left(w\, r^2\rho\chi\right)}{\partial \eta} \\ = \frac{1}{h_\eta}\frac{\partial}{\partial \eta}\left(\frac{\rho\, d_\eta}{h_\eta}\frac{\partial \chi}{\partial \eta}\right) + P(\rho\chi) + F(\rho\chi).\end{aligned} \tag{2.10}$$

Horizontal diffusion is neglected here and d_η is the vertical diffusion coefficient. The velocity components in the λ, ϕ and η direction are respectively u, v and w, and $h_\eta = |\partial R/\partial \eta|$ is the scale factor from the coordinate transformation. Some existing models neglect the variation in r and therefore they replace the radial distance r with the (constant) radius of the earth $a = 6378$ km. Then both factors r^2 cancel in the part of Equations (2.9) and (2.10) describing vertical advection.

However, the earth is not a perfect sphere because there are mountains. This complicates the coordinate transformation, because the hybrid coordinate system is a nonorthogonal coordinate system, see Figure 2.2. The complete, general transformation is given in [BR97], where the authors show that the contravariant velocity component should be used. Equation (2.10) can still be used in global models, because even on a relatively fine grid ($1° \times 1°$) the error due to the simplified coordinate transformation is small compared to other model errors. For regional models using much finer grids this is no longer true.

2.3.3 Mass fluxes

In the finite volume approach, Equation (2.10) is interpreted in terms of in- and outgoing fluxes. Existing models such as TM3, multiply ρ with the volume V to obtain

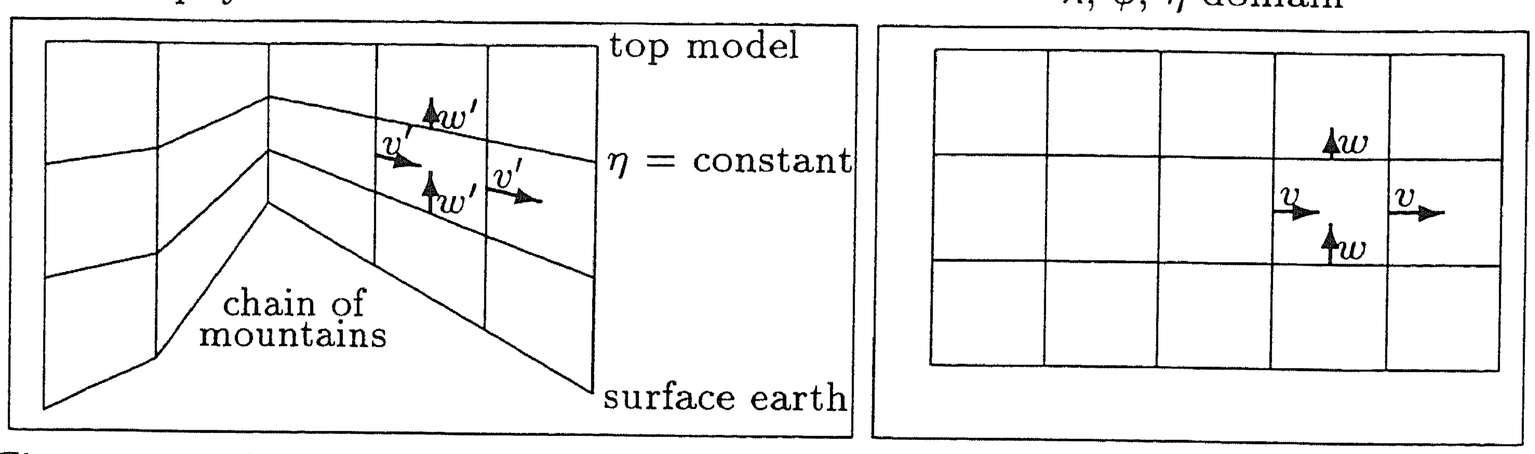

Figure 2.2: Orthogonal velocity components in the η-coordinate system are not orthogonal in the physical domain.

particles per grid cell in Equation (2.10) and air mass m per grid cell [kg] in Equation (2.9). The velocities u, v and w are multiplied with density times area of the cell interface to obtain mass fluxes $\mathcal{F}$. This approach has at least two advantages. Firstly, the calculation of h_η can be performed during the preprocessing of the meteorological input, and secondly, the mass flux is well defined, even in a nonorthogonal coordinate system. The mass flux of a cell interface is the total amount of air [in kg] that flows through that particular cell interface in 1 second. Observe that the discretized advection part of Equation (2.10) can be written as a rate of change in mass due to inflow and outflow of (clean) air m with tracers. Diffusion and sub-grid scale parametrizations (described in Section 6.2.2), lead to fluxes f describing vertical transport by means of sub-grid scale vertical transport processes, corresponding to P in Equation (2.10). This leads to an ODE for all grid cells in terms of $m = \rho V$.

$$\begin{aligned}\frac{\mathrm{d}}{\mathrm{d}t} m_{i,j,k} = {} & \mathcal{F}_{\lambda_{i-\frac{1}{2},j,k}} - \mathcal{F}_{\lambda_{i+\frac{1}{2},j,k}} + \mathcal{F}_{\phi_{i,j-\frac{1}{2},k}} - \mathcal{F}_{\phi_{i,j+\frac{1}{2},k}} \\ & + \mathcal{F}_{\eta_{i,j,k-\frac{1}{2}}} - \mathcal{F}_{\eta_{i,j,k+\frac{1}{2}}}\end{aligned} \tag{2.11}$$

$$\begin{aligned}\frac{\mathrm{d}}{\mathrm{d}t} m_{i,j,k}\, \chi_{i,j,k,l} = {} & \mathcal{F}_{\lambda_{i-\frac{1}{2},j,k}}\, \chi_{i-\frac{1}{2},j,k,l} - \mathcal{F}_{\lambda_{i+\frac{1}{2},j,k}}\, \chi_{i+\frac{1}{2},j,k,l} \\ & + \mathcal{F}_{\phi_{i,j-\frac{1}{2},k}}\, \chi_{i,j-\frac{1}{2},k,l} - \mathcal{F}_{\phi_{i,j+\frac{1}{2},k}}\, \chi_{i,j+\frac{1}{2},k,l} \\ & + \mathcal{F}_{\eta_{i,j,k-\frac{1}{2}}}\, \chi_{i,j,k-\frac{1}{2},l} - \mathcal{F}_{\eta_{i,j,k+\frac{1}{2}}}\, \chi_{i,j,k+\frac{1}{2},l} \\ & + \textstyle\sum_{n=1}^{N_\eta} m_{i,j,n} \chi_{i,j,n,l} \left(f_{i,j,k-\frac{1}{2},n} - f_{i,j,k+\frac{1}{2},n} \right) \\ & + F_l \left(\rho_{i,j,k}\, \chi_{i,j,k,1 \ \ldots \ N_c} \right).\end{aligned} \tag{2.12}$$

The air mass fluxes [kg s^{-1}] through cell interfaces are $\mathcal{F}_\lambda$, $\mathcal{F}_\phi$ and $\mathcal{F}_\eta$; N_η is the number of vertical layers and $m_{i,j,k}$ is the air mass in a grid point and $\chi_{i,j,k,l}$ is the mixing ratio of species l in grid point (i,j,k). In Equation (2.12) both the ratio tracer mass divided by air mass and tracer concentration divided by the density of air can be used.

```
#EQUATIONS {A small example.}
 OH + CO {+ O2} = CO2 + HO2 : 1.5E-13*(1.+6E-6*press) ;
OH + CH4 {+ O2} = H2O + CH3O2 : 2.36E-12*exp(-1710./TEMP) ;
O3P + O2 = O3 : 1.5E-11 ;
NO + O3 = NO2 + O2 : 1.8E-12 * exp(-1370. / TEMP) ;
NO + HO2 = NO2 + OH : 3.7E-12 * exp(240. / TEMP) ;
NO2 + O3 = NO3 + O2 : 1.2E-13 * exp(-2450. / TEMP);
NO2 + NO3 = NO + NO2 + O2 : 2.3E-13 * exp(-1E3 / TEMP) ;
NO + NO3 = 2 NO2 : 1.3E-11 * exp(250. / TEMP) ;
O3 + hv = O3P + O2 : 1.23E-3 * exp(-.6 * sec_Z) ;
O3 + hv = O1D + O2 : 2E-4 * exp(-1.4 * sec_Z) ;
NO2 + hv = NO + O3P : 1.45E-2 * exp(-.4 * sec_Z) ;
NO3 + hv = NO + O2 : 3.53E-2 * exp(-8.1E-2 * sec_Z) ;
NO3 + hv = NO2 + O3P : 8.94E-2 * exp(-5.92E-2 * sec_Z) ;
```

Figure 2.3: An example of a KPP input file.

In Chapter 6, where we use realistic meteorological data, we solve Equation (2.12). In the remaining chapters we solve Equation (2.10).

2.4 Chemistry

An important development with respect to the chemistry part of the model is KPP, a chemical Kinetics PreProcessor [DIS95]. This program reads an input file with reactions and generates `C` and `Fortran 77` code for the function $F(c)$, the Jacobian $J(c) = \partial F(c) / \partial c$, the LU decomposition of the Jacobian, the backsolve $LUx = y$, and the matrix-vector multiplication $J(c)\,c$. An example of a KPP input file is found in Figure 2.3. The function $F(c)$ is given in aggregate form ($\dot{c} = F(c)$) and splitted in production and loss terms ($\dot{c} = P(c) - L(c)\,c$). KPP generates both sparse matrix code and full matrix code. Only the sparse LU decomposition is coded using indirect addressing. The sparsity routines have turned out to be very efficient in reducing the numerical linear algebra overhead [SPCD96].

The current version of KPP generates code for a box-model only. For a 3D model, you have to add grid indices by hand, which takes hand-work. Also the implementation of the Gauss-Seidel iteration used in Twostep [Ver94, VS95] takes hand-work, even for the box-model.

However, a big advantage of using KPP is that it is very easy to study several chemistry models. While developing numerical solvers it is important to measure stability and accuracy versus cpu time on a large set of different chemistry models, otherwise you might end up with a numerical solver which works very well on a particular chemistry used while developing your solver, but which might fail on a new chemistry model.

chemistry model	*LU* indirect		*LU* direct	
	LU	total	*LU*	total
RIVM	20	126	10	117
CBM-IV	56	292	41	286
WET	120	729	100	726

Table 2.2: cpu times on a workstation [s] for *LU* decomposition and total cpu time for three chemistry models. Left with indirect addressing and right with direct addressing.

We benchmarked our solvers on state of the art chemistry models, such as EMEP [Sim93], CBM-IV [GWKD89] and WET [Mat95], and on models which are derived from these.

2.4.1 *LU* decomposition in KPP

Recently, we found out how to write a subroutine which generates a new subroutine that also codes the sparse *LU* decomposition using direct addressing, similar as the routine for the backsolve. This leads to a better performance on both workstations and vector processors. To avoid any confusion, we state that the timings presented in this thesis and in e.g. [SVB+97] are obtained with the sparse *LU* decomposition using indirect addressing.

To see the impact of this improvement, we repeated the boxmodel tests of Chapter 7. To get statistical significant differences in cpu time, we extended the test from 14 days to 140 days. In Table 2.2 cpu times are given for three chemistry models, integrated with ROS2 using a fixed time step of 10 minutes.

We observe that the *LU* decomposition consumes roughly 17% of the total cpu time, and the *LU* decomposition with direct addressing is 17-50% faster than indirect addressing, depending on the chemistry used.

Chapter 3

An efficient horizontal advection scheme for the modeling of global transport of constituents

Abstract

In this chapter we consider a dimensional-splitting scheme for horizontal advection on a sphere with a uniform longitude-latitude grid. The 1D subprocesses that arise within the splitting are solved with an explicit finite-volume type scheme, which is made unconditionally stable by allowing the stencil to vary with the Courant numbers. The scheme is made positive by flux-limiting. For the inaccuracies at the poles some special measures are discussed. Numerical tests show that the scheme is almost shape preserving and conservative, and it gives accurate results at low computational costs.

Note: This chapter has been published in *Monthly Weather Review*, **123**, 1995, pages 3554 - 3564.

3.1 Introduction

Global atmospheric transport models are becoming increasingly important for the analysis and description of the distribution of trace gases and particles. Climate change is closely related to the global tropospheric distribution of trace constituents. These originate mostly from regional or local sources mainly concentrated over land,

and, consequently, in the Northern Hemisphere. Once released, a complicated atmospheric process, involving transport, chemistry and depositions, takes place, which, in view of the atmospheric lifetime of many species and typical horizontal transport velocities, involves the entire global troposphere.

Current modeling of global tracer transport can be done on-line and off-line. In the latter case the tracer is advected by a wind field which originates from a circulation or weather forecast model, such as the European Centre for Medium Range Weather Forecasts (ECMWF) model. Examples of these models are the GISS model [JPWM87], from which the TM2 model [HK89] was derived, and the Moguntia model [Zim87]. More recently, on-line modeling is developing where the tracer advection equation is part of a complete atmospheric circulation model (e.g., [FRSW92]). This has the obvious advantage that a consistent set of equations is used with the possibility of interaction between the tracer and the (thermo)dynamics. In both approaches the horizontal advection is an important issue.

For the numerical solution of horizontal advection in Cartesian coordinates numerous schemes exist (see, e.g., [Roo87, VK93]). The application to global advection offers some additional problems that are associated with the spherical geometry. These aspects in particular will be emphasized here.

The mathematical equation describing the horizontal advection on a sphere for a concentration, or vector of concentrations, $c(t, \lambda, \phi)$ is given by

$$\frac{\partial c}{\partial t} + \frac{1}{a\cos(\phi)}\left(\frac{\partial (uc)}{\partial \lambda} + \frac{\partial (\cos(\phi)vc)}{\partial \phi}\right) = 0, \tag{3.1}$$

where $\lambda \in [0, 2\pi)$ and $\phi \in [-\pi/2, \pi/2]$ are the longitude and latitude coordinates, a is the radius of the sphere, and u, v the given wind velocities (m s^{-1}) in the longitudinal and latitudinal direction, respectively. The advection equation (3.1) becomes singular at the poles $\phi = \pm\,\pi/2$, due to the underlying mapping of the sphere on the λ, ϕ-plane. This may cause inaccuracies in numerical schemes as well as stability problems. A thorough review with many references can be found in [Wil92].

In the numerical transport of trace species it is important to avoid undershoot, since this can result in unrealistic negative concentrations. If the species is chemically active this may lead to fatal numerical problems in the treatment of chemical reactions. Overshoot can also be undesirable. For example, in the case of transport of water vapour it will lead to unwanted condensation and cloud formation.

In this paper we consider an advection scheme based on dimensional splitting. The underlying grid is assumed to be a uniform longitude-latitude grid, which gives a very simple data structure. The 1D subprocesses that arise within the splitting are solved with an explicit third-order finite-volume type scheme, used in [HT94] for planar advection. Positivity is obtained by flux-limiting. The resulting scheme is not completely monotone; overshoot and undershoot (against a positive background concentration) is still possible, but numerical tests show that the magnitude of these over- and undershoots are small.

The scheme is made unconditionally stable by allowing the stencil (numerical domain of dependence) to vary with the local Courant number. In this way, small time steps near the poles can be avoided.

However, due to the large Courant numbers that will arise, accuracy near the poles remains a problem with this approach. To suppress these inaccuracies we shall propose two modifications, namely mixing inside polar caps and deformation corrections (where the scheme is corrected such that background concentrations are not deformed). With both modifications the results are satisfactory. The resulting schemes have good shape preserving properties, maintain positivity, and appear to be very efficient in terms of computational costs versus accuracy. Comparisons will be made with the semi-Lagrangian methods of Smolarkiewicz and Rasch [SR91] for a solid body rotation test on the sphere.

In Section 3.2 some results of Hundsdorfer and Trompert [HT94] concerning advection on a plane are reviewed and the extension to unconditional stability is discussed. Section 3.3 contains the tests for advection on the sphere, with modifications at the poles and numerical comparisons.

3.2 Planar advection

3.2.1 Unconditional stability for 1D advection

For the 1D advection equation

$$c_t + (uc)_x = 0, \tag{3.2}$$

with subscripts t, x denoting partial derivatives, we consider numerical schemes in conservation form

$$c_i^{n+1} = c_i^n + \left(F^n_{i-\frac{1}{2}} - F^n_{i+\frac{1}{2}} \right). \tag{3.3}$$

Here c_i^n approximates the average value of $c(t,x)$ over the cell $\Omega_i = (x_{i-\frac{1}{2}}, x_{i+\frac{1}{2}})$ at time $t_n = n\Delta t$, and $F^n_{i-\frac{1}{2}}, F^n_{i+\frac{1}{2}}$ give the inflow and outflow at the cell boundaries. The grid will be assumed to be equidistant, with cell centers $x_i = i\Delta x$, but generalizations to arbitrary grids are easily obtained. In the following the superscripts n are omitted.

For the computation of the fluxes $F_{i+\frac{1}{2}}$ we shall allow the stencil to vary with the Courant number. In this way it is easy to achieve unconditional stability while maintaining explicitness of the scheme. For example, the first-order upwind, donor cell scheme, is given by

$$F_{i+\frac{1}{2}} = \begin{cases} \nu_{i+\frac{1}{2}} c_i, & u_{i+\frac{1}{2}} \geq 0, \\ -\nu_{i+\frac{1}{2}} c_{i+1}, & u_{i+\frac{1}{2}} < 0, \end{cases} \tag{3.4}$$

with $\nu_{i+\frac{1}{2}} = \left| u_{i+\frac{1}{2}} \right| \Delta t/\Delta x$ the local Courant number at the cell boundary $x_{i+\frac{1}{2}}$. In this form the scheme is stable under the CFL restriction $\max_i \nu_{i+\frac{1}{2}} \leq 1$. The stability

restriction is avoided by taking

$$F_{i+\frac{1}{2}} = \begin{cases} \tilde{\nu}_{i+\frac{1}{2}} c_{i-k} + (c_{i-k+1} + \cdots + c_i), & u_{i+\frac{1}{2}} \geq 0, \\ -\tilde{\nu}_{i+\frac{1}{2}} c_{i+k+1} - (c_{i+1} + \ldots + c_{i+k}), & u_{i+\frac{1}{2}} < 0, \end{cases} \tag{3.5}$$

where

$$k = k_{i+\frac{1}{2}} = \lfloor \nu_{i+\frac{1}{2}} \rfloor \quad \text{and} \quad \tilde{\nu}_{i+\frac{1}{2}} = \nu_{i+\frac{1}{2}} - k_{i+\frac{1}{2}}, \tag{3.6}$$

and $\lfloor \nu \rfloor$ denotes the largest integer $\leq \nu$. Inserting this into (3.3), we see that for constant u the same formula as for Courant numbers ≤ 1 is applied, together with a shift over k grid points. Therefore, the scheme will be unconditionally stable (for constant u). This is as with semi-Lagrangian methods, but due to the conservation form of the scheme we maintain the mass-conservation property. We note that a scheme similar to (3.5) was introduced by LeVeque [LeV82] for nonlinear conservation laws; a somewhat related technique can be found in [Bre84].

Instead of first-order upwind we shall use an upwind-biased scheme that is third-order accurate for constant velocities. Let

$$d_0(\nu) = \tfrac{1}{6}(2-\nu)(1-\nu), \quad d_1(\nu) = \tfrac{1}{6}\left(1-\nu^2\right). \tag{3.7}$$

If $u_{i+\frac{1}{2}} \geq 0$ we set

$$\begin{aligned} F_{i+\frac{1}{2}} \quad &= \tilde{\nu}_{i+\frac{1}{2}} \left(c_{i-k} \; + \; d_0\left(\tilde{\nu}_{i+\frac{1}{2}}\right)(c_{i+1-k} - c_{i-k}) \right. \\ &\left. + \; d_1\left(\tilde{\nu}_{i+\frac{1}{2}}\right)(c_{i-k} - c_{i-1-k}) \right) \; + \; \sum_{l=i-k+1}^{i} c_l, \end{aligned} \tag{3.8}$$

whereas for $u_{i+\frac{1}{2}} < 0$,

$$\begin{aligned} F_{i+\frac{1}{2}} \quad &= -\tilde{\nu}_{i+\frac{1}{2}} \left(c_{i+k+1} \; + \; d_0\left(\tilde{\nu}_{i+\frac{1}{2}}\right)(c_{i+k} - c_{i+k+1}) \right. \\ &\left. + \; d_1\left(\tilde{\nu}_{i+\frac{1}{2}}\right)(c_{i+k+1} - c_{i+k+2}) \right) \; - \; \sum_{l=i+1}^{i+k} c_l, \end{aligned} \tag{3.9}$$

with $k = k_{i+\frac{1}{2}}$ and $\tilde{\nu}_{i+\frac{1}{2}}$ given by (3.6).

If u is constant and the Courant number is bounded by 1 the scheme reduces to the third-order scheme which uses a four-point upwind biased stencil. It was shown by Strang [Str62] that schemes of this optimal-order type are stable up to Courant number equal to 1. Unconditional stability, for constant u, of the above formula follows again by observing that we apply in fact the same formula as for Courant number less than or equal to 1, only on a shifted stencil.

To achieve positivity we will use a Sweby-type flux limiter (See [Swe84, Zal87, LeV92, Kor93]). Detailed motivation for the particular limiter used here can be found in [HT94]. The limiter is based on the ratios

$$\theta_i = \frac{c_i - c_{i-1}}{c_{i+1} - c_i} \tag{3.10}$$

and the limiter function is

$$\psi(\nu,\theta) = \max\left(0, \min\left(1, d_0(\nu) + d_1(\nu)\theta, \frac{1-\nu}{\nu}\theta\right)\right). \tag{3.11}$$

In actual computations a small number, for example, 10^{-10}, is added to the denominators of θ and $(1-\nu)/\nu$ to avoid division by zero. The limited scheme reads as follows: if $u_{i+\frac{1}{2}} > 0$ then $F_{i+\frac{1}{2}}$ is computed from

$$F_{i+\frac{1}{2}} = \tilde{\nu}_{i+\frac{1}{2}}\left(c_{i-k} + \psi\left(\tilde{\nu}_{i+\frac{1}{2}}, \theta_{i-k}\right)\left(c_{i+1-k} - c_{i-k}\right)\right) + \sum_{l=i-k+1}^{i} c_l, \tag{3.12}$$

with $k = k_{i+\frac{1}{2}}$ and $\tilde{\nu}_{i+\frac{1}{2}}$ as before, and if $u_{i+\frac{1}{2}} < 0$ we use

$$F_{i+\frac{1}{2}} = -\tilde{\nu}_{i+\frac{1}{2}}\left(c_{i+k+1} + \psi\left(\tilde{\nu}_{i+\frac{1}{2}}, \frac{1}{\theta_{i+k+1}}\right)\left(c_{i+k} - c_{i+k+1}\right)\right) - \sum_{l=i+1}^{i+k} c_l. \tag{3.13}$$

Note that with $\psi(\nu,\theta) = d_0(\nu) + d_1(\nu)\theta$ the original, unlimited scheme would have been reobtained.

The above formulas depend, through $\nu_{i+\frac{1}{2}}$, on the velocity $u_{i+\frac{1}{2}}$. From a physical interpretation of the flux form (3.3), in terms of inflow and outflow for the cells, it follows that $u_{i+\frac{1}{2}}$ should be taken as $\left(x_{i+\frac{1}{2}} - \bar{x}_{i+\frac{1}{2}}\right)/\Delta t$, where $\bar{x}_{i+\frac{1}{2}}$ is the departure point at time t_n of the characteristic $x'(t) = u(t, x(t))$ that passes through $x_{i+\frac{1}{2}}$ at time t_{n+1} (see [HT94], for instance). Therefore, if we take simply $u_{i+\frac{1}{2}} = u\left(x_{i+\frac{1}{2}}\right)$, the order will only be 1 for non-constant velocities. In this paper we will apply the second-order formula

$$u_{i+\frac{1}{2}} = \left(u - \tfrac{1}{2}\Delta t\, u_x u\right)\left(x_{i+\frac{1}{2}}\right). \tag{3.14}$$

The derivatives u_x appearing here are computed by second-order difference quotients.

Higher order approximations for the departure point could be easily derived, as well as modifications in case u also varies in time, see [HT94]. However, we shall apply the 1D scheme within a second-order splitting setting, and in our applications the wind fields will be held constant for fairly long time intervals.

3.2.2 Second-order splittings

To solve the 2D advection equation on a plane

$$c_t + (uc)_x + (vc)_y = 0, \tag{3.15}$$

we consider dimensional splitting. In it simplest form this reads as follows: if $c_n(x,y)$ is an approximation to the exact solution at time t_n, we solve subsequently

$$\begin{aligned} c^*_t + (uc^*)_x &= 0 \quad (t_n \le t \le t_{n+1}), \\ c^{**}_t + (vc^{**})_y &= 0 \quad (t_n \le t \le t_{n+1}), \end{aligned} \tag{3.16}$$

with input values $c^*(t_n, x, y) = c_n(x, y)$ and $c^{**}(t_n, x, y) = c^*(t_{n+1}, x, y)$, giving a new approximation $c_{n+1}(x, y) = c^{**}(t_{n+1}, x, y)$ at time level t_{n+1}.

It is well known that this type of splitting will introduce a first-order temporal error. In fact, by a simple analysis (see [HT94]) it follows that the splitting procedure will solve the *modified* equation

$$c_t + (\underline{u}c)_x + (\underline{v}c)_y = \mathcal{O}\left(\Delta t^2\right) \tag{3.17}$$

with perturbed velocities

$$\underline{u} = u - \tfrac{1}{2}\Delta t\, u_y v, \qquad \underline{v} = v + \tfrac{1}{2}\Delta t\, u v_x. \tag{3.18}$$

It follows that second-order splitting is achieved by using in the splitting procedure not the original velocities u and v, but instead

$$\overline{u} = u + \tfrac{1}{2}\Delta t\, u_y v, \qquad \overline{v} = v - \tfrac{1}{2}\Delta t\, u v_x. \tag{3.19}$$

This precisely counteracts the first-order perturbations in the above modified equation.

As observed by Bott [Bot92] first-order splitting may give qualitatively bad results with deformational flow fields. We have repeated the experiment of Bott [Bot92] [on the deformation test of Smolarkiewicz [Smo82] with background]: in the flow field described by the stream function

$$\varphi(x, y) = 8 \sin\left(\frac{\pi x}{25}\right) \cos\left(\frac{\pi y}{25}\right), \tag{3.20}$$

for $0 \le x, y \le 100$, with $u = -\varphi_y$ and $v = \varphi_x$, a cone with center $(x, y) = (50, 50)$, base radius 15, and height 4 is placed on top of a background concentration equal to 1. In this flow, the cone should be strongly deformed but the background should remain constant since the velocity field is divergence-free.

In Figure 3.1 the numerical solutions are displayed for first-order and second-order splitting, with $\Delta t = 0.7$ and $\Delta x = \Delta y = 1$ after 38 time steps. The fractional steps are solved by the limited scheme discussed in the previous section. With first-order splitting we use the velocities at cell interfaces according to (3.14); that is

$$\begin{aligned} u_{i\pm\frac{1}{2},j} &= \left(u - \tfrac{1}{2}\Delta t\, u_x u\right)(x_{i\pm\frac{1}{2}}, y_j), \\ v_{i,j\pm\frac{1}{2}} &= \left(v - \tfrac{1}{2}\Delta t\, v_y v\right)(x_i, y_{j\pm\frac{1}{2}}). \end{aligned} \tag{3.21}$$

For the second-order splitting we also take into account the modification (3.19), which leads to

$$\begin{aligned} u_{i\pm\frac{1}{2},j} &= \left(u - \tfrac{1}{2}\Delta t\, u_x u + \tfrac{1}{2}\Delta t\, u_y v\right)\left(x_{i\pm\frac{1}{2}}, y_j\right), \\ v_{i,j\pm\frac{1}{2}} &= \left(v - \tfrac{1}{2}\Delta t\, v_y v - \tfrac{1}{2}\Delta t\, u v_x\right)\left(x_i, y_{j\pm\frac{1}{2}}\right). \end{aligned} \tag{3.22}$$

It is clear from Figure 3.1 that with first-order splitting not only the cone but also the background concentration is strongly deformed (without the modification (3.14) it

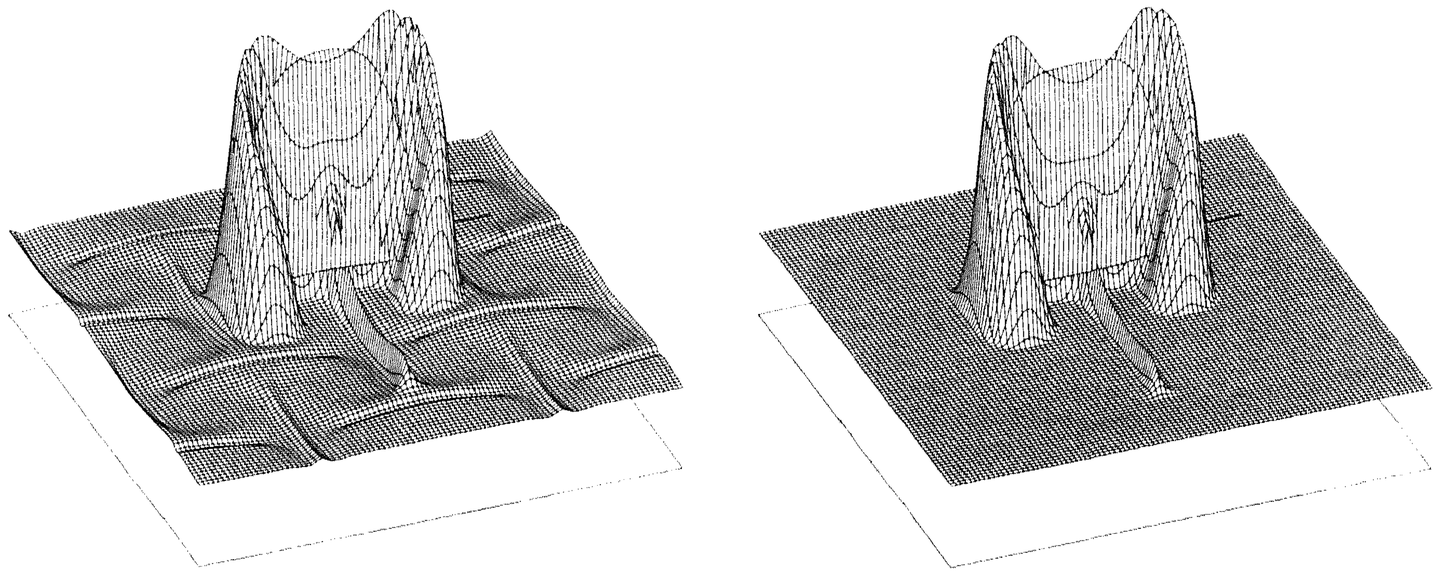

Figure 3.1: First-order (left) and second-order splitting (right) in deformation test.

is even worse, similar to Figure 9b in [Bot92]). This deformation is caused by the fact that the modified equation is no longer divergence-free. The result with second-order splitting is satisfactory. Although there still is a slight deformation of the background concentration, this is hardly visible anymore.

Remark. Second-order accuracy is also achieved with the so-called Strang splitting [Str68, LeV92],

$$\begin{aligned} c_t^* + (uc^*)_x &= 0 \quad \left(t_n \le t \le t_{n+\frac{1}{2}}\right), \\ c_t^{**} + (vc^{**})_y &= 0 \quad \left(t_n \le t \le t_{n+1}\right), \\ c_t^{***} + (uc^{***})_x &= 0 \quad \left(t_{n+\frac{1}{2}} \le t \le t_{n+1}\right), \end{aligned} \tag{3.23}$$

with input values $c^*(t_n, x, y) = c_n(x, y)$, $c^{**}(t_n, x, y) = c^*\left(t_{n+\frac{1}{2}}, x, y\right)$ and $c^{***}\left(t_{n+\frac{1}{2}}, x, y\right) = c^{**}(t_{n+1}, x, y)$. Now $c^{***}(t_{n+1}, x, y)$ gives the new approximation at time level t_{n+1}. In some tests we found little difference between this form of splitting and the form (3.16) with modified wind field (3.22). If at each time level t_n output is demanded, then Strang splitting will be more expensive. In the following, we therefore consider splitting with the modified wind field (3.22).

3.3 Advection on the sphere

3.3.1 The splitting scheme

Consider the advection equation (3.1) on a sphere. In the following we denote, for brevity, $\Gamma = \cos(\phi)$. The dimension of the velocities u, v is meter per second. The

relevant velocities on the λ, ϕ-plane are

$$U = \frac{u}{a\Gamma} \quad \text{and} \quad V = \frac{v}{a}. \tag{3.24}$$

Note that U becomes very large near the poles, reflecting the fact that the mesh width on the sphere in the λ direction becomes small.

Equation (3.1) can be written as

$$\frac{\partial c}{\partial t} + \frac{\partial (Uc)}{\partial \lambda} + \frac{1}{\Gamma} \frac{\partial (\Gamma V c)}{\partial \phi} = 0. \tag{3.25}$$

We consider the splitting scheme

$$c^*_{i,j} = c^n_{i,j} + \left(F^n_{i-\frac{1}{2},j} - F^n_{i+\frac{1}{2},j} \right), \tag{3.26}$$

$$c^{n+1}_{i,j} = c^*_{i,j} + \frac{1}{\Gamma_j} \left(\Gamma_{j-\frac{1}{2}} F^*_{i,j-\frac{1}{2}} - \Gamma_{j+\frac{1}{2}} F^*_{i,j+\frac{1}{2}} \right), \tag{3.27}$$

where the 1D fluxes are computed as in Section 3.2, using Courant numbers

$$\begin{aligned} \nu_{i\pm\frac{1}{2},j} &= \frac{\Delta t}{\Delta \lambda} \left| U_{i\pm\frac{1}{2},j} \right|, \\ \nu_{i,j\pm\frac{1}{2}} &= \frac{\Delta t}{\Delta \phi} \left| V_{i,j\pm\frac{1}{2}} \right|. \end{aligned} \tag{3.28}$$

(In fact, the limiter in the ϕ direction will be slightly modified to incorporate the different factors Γ in (3.27), see the Appendix 3.4.1.)

The computations are performed on a uniform $2m \times m$ grid with cell centers (λ_i, ϕ_j) given by $\lambda_i = (i - \frac{1}{2})\Delta\lambda$ and $\phi_j = (j - \frac{1}{2})\Delta\phi - \frac{1}{2}\pi$ for $i = 1, \cdots, 2m$, $j = 1, \cdots, m$ and $\Delta\lambda = \Delta\phi = \pi/m$. The cell boundaries are located at $\lambda_{i+\frac{1}{2}} = \frac{1}{2}(\lambda_i + \lambda_{i+1})$ and $\phi_{j+\frac{1}{2}} = \frac{1}{2}(\phi_j + \phi_{j+1})$. At the poles we have a no-flux condition, that is, $\Gamma_{j-\frac{1}{2}} F^*_{i,j-\frac{1}{2}} = 0$ for $j = 1$ and $\Gamma_{j+\frac{1}{2}} F^*_{i,j+\frac{1}{2}} = 0$ for $j = m$, reflecting the fact that the physical length of these cell boundaries is zero. Furthermore, for cells that are located downstream of the poles, the calculation of the fluxes in ϕ direction may need values $c^*_{i,j}$ across the poles. These are found by constant extrapolation,

$$c^*_{i,0} = c^*_{i,1}, \quad c^*_{i,m+1} = c^*_{i,m}. \tag{3.29}$$

(Interpolation, using values across the poles, gave worse results, due to the fact that the $c^*_{i,j}$ are nonphysical quantities that are discontinuous across the poles.)

The modified velocities at the cell boundaries will be taken as

$$U_{i\pm\frac{1}{2},j} = \left(U - \tfrac{1}{2}\Delta t \frac{\partial U}{\partial \lambda} U + \tfrac{1}{2}\Delta t \frac{\partial U}{\partial \phi} V \right) \left(\lambda_{i\pm\frac{1}{2}}, \phi_j \right), \tag{3.30}$$

$$V_{i,j\pm\frac{1}{2}} = \left(V - \tfrac{1}{2}\Delta t \, \frac{1}{\Gamma}\frac{\partial \Gamma V}{\partial \phi} V - \tfrac{1}{2}\Delta t \, U \frac{\partial V}{\partial \lambda} \right) \left(\lambda_i, \phi_{j\pm\frac{1}{2}} \right). \tag{3.31}$$

These velocities are derived in the same way as for advection on a plane [see formula (3.22)], with the aim of achieving second order accuracy in time. However, the derivation is based on Taylor expansions of the error terms. Near the poles these Taylor expansions loose their significance, since the expansions will involve powers of $\Delta t/\Gamma$ instead of genuine powers of Δt. At the grid points nearest to the pole we have $\Gamma \sim \Delta\phi$, so that the asymptotic considerations are only valid if $\Delta t/\Delta\phi \to 0$.

As we shall see, the scheme is indeed inaccurate near the poles if both Δt and $\Delta\phi$ are small. In fact, the local errors near the poles are only slightly smaller than with first-order splitting. Nevertheless, second-order is preferable for accuracy outside the polar regions for deformational flows as in Figure 3.1. For the errors at the poles some special measures will be considered.

Remark. It is clear from the conservation form of (3.26), (3.27) that the scheme will conserve the quantity

$$\Delta\lambda \, \Delta\phi \sum_{i=1}^{2m} \sum_{j=1}^{m} \Gamma_j \, c_{i,j}, \tag{3.32}$$

which approximates the total mass of the concentration field c on the sphere.

3.3.2 Solid-body rotation test on the sphere

The reduction of accuracy near the poles is well illustrated with the solid-body rotation test of Williamson and Rasch [WR89]. As in [SR91], we consider a cone shaped initial profile given by

$$c_0(\lambda, \phi) = \max\left(0, 1 - r(\lambda, \phi)/R\right), \tag{3.33}$$

where

$$r(\lambda, \phi) = 2\sqrt{\left(\cos(\phi)\sin\left(\tfrac{1}{2}\left(\lambda - \tfrac{3}{2}\pi\right)\right)\right)^2 + \left(\sin\left(\tfrac{1}{2}\phi\right)\right)^2} \tag{3.34}$$

and $R = 7\pi/m$. Here m is the number of grid points in ϕ direction, taken as $m = 64$, and R is the base radius of the cone corresponding to seven grid points. The center of the cone is $(\frac{3}{2}\pi, 0)$. The velocities in this test are

$$\begin{aligned} u(\lambda, \phi) &= 2\pi(\cos(\beta)\cos(\phi) + \sin(\beta)\cos(\lambda)\sin(\phi)), \\ v(\lambda, \phi) &= -2\pi\sin(\beta)\sin(\lambda), \end{aligned} \tag{3.35}$$

with angle

$$\beta = \pi/2. \tag{3.36}$$

With this velocity field the cone profile will have completed at time $t = 1$ one full rotation with trajectory over both poles. This will be covered in 256 time steps (maximal Courant numbers of about 20) with the flux-limited splitting scheme.

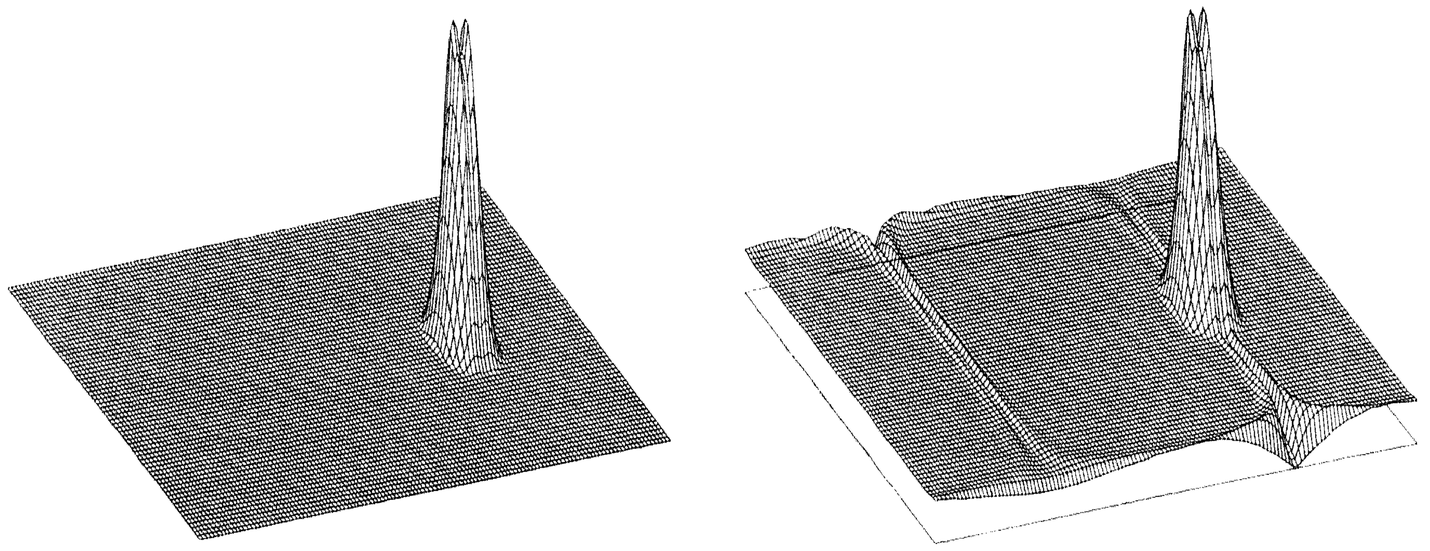

Figure 3.2: Cone tests without background (left) and with background (right).

The left picture of Figure 3.2 shows a 3D plot of the numerical result on the $\lambda - \phi$ plane after one rotation. We see a strong deformation at the center of the cone, which has crossed both poles.

To gain more insight, it is useful to consider a background concentration. In the right picture of Figure 3.2 the result is shown for the initial profile

$$\tilde{c}_0(\lambda, \phi) = 0.1 \; + \; 0.9 \; c_0(\lambda, \phi). \tag{3.37}$$

We now see that the background also is strongly deformed as it flows over the poles. In each step an $\mathcal{O}(1)$ error is introduced at the poles, and these errors are advected in the direction of the equator along the streamlines. The errors become less prominent during the advection towards the equator due to the fact that the area of the cells (on the sphere) become larger.

The errors at the poles are caused by the splitting itself as well as by the approximations for the 1D processes. These errors can be made small by taking a sufficiently small time step, say 1/5120 with maximal Courant number of about 1, but this implies of course a large increase in computational work. In the following sections we will discuss some special measures to suppress the large errors near the poles while maintaining the computational efficiency.

3.3.3 Polar caps and polar mixing

Allen et al. [ADRG91] considered a splitting scheme where the 1D processes were subject to a CFL restriction. In order to relax this restriction they introduced polar caps, considered as cells with homogeneous distribution. With such polar caps the splitting scheme is applied farther away from the poles. This will lead to smaller Courant numbers, and also the splitting errors will become less severe. On the other

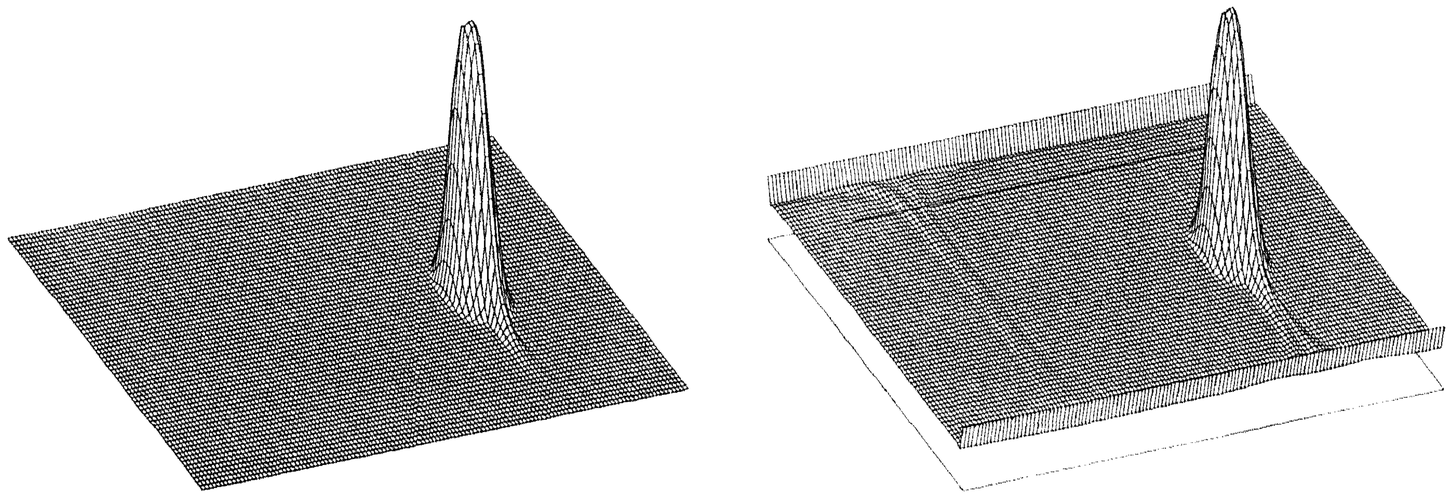

Figure 3.3: Cone tests for polar caps.

hand, the polar caps will introduce diffusion for profiles passing the poles, and so these caps should not be taken to be too large. We note that the 1D scheme that was used in [ADRG91], based on [vL77], is already somewhat more diffusive than the present scheme (3.11)-(3.13).

In Figure 3.3 the numerical solutions are shown for polar caps with radius $\Delta\phi$ (the first and last row in the plots correspond to values inside the two caps). The diffusive effect of the polar caps is clearly seen for the profile without background. The picture with background shows that also the splitting error is still quite large with this time step size. We note that in [ADRG91] smaller time steps were taken near the poles, but in their Figure 4 a similar diffusive effect can be observed.

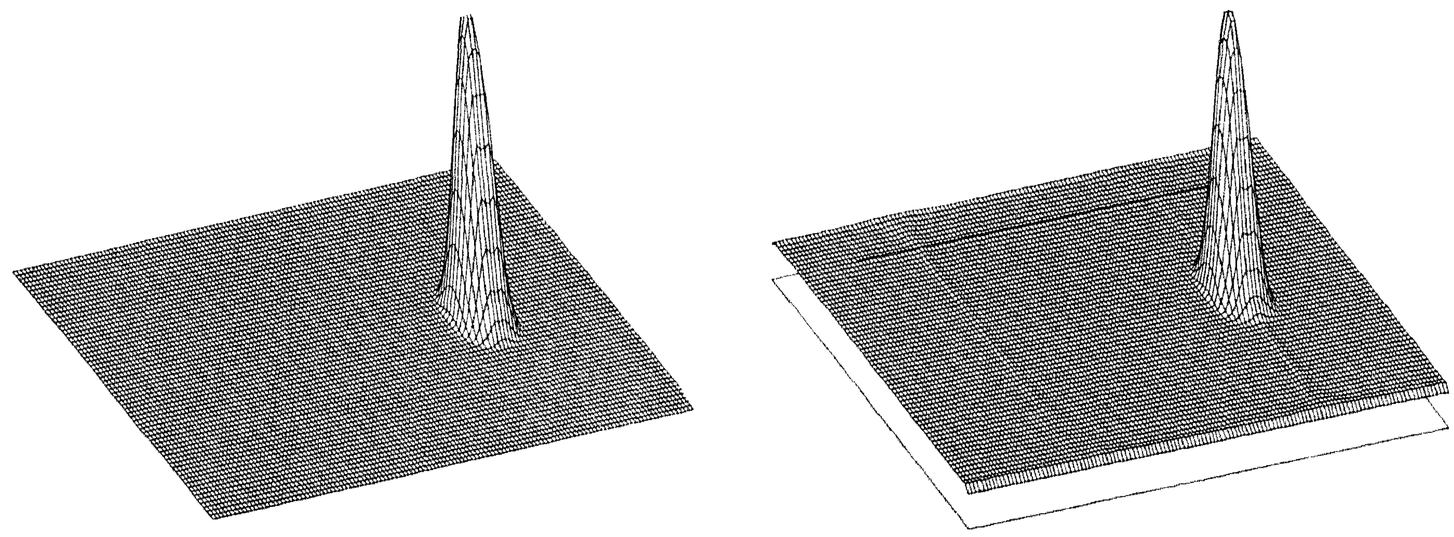

Figure 3.4: Cone tests for polar mixing.

By experiment we found that better results are obtained if the scheme is simply

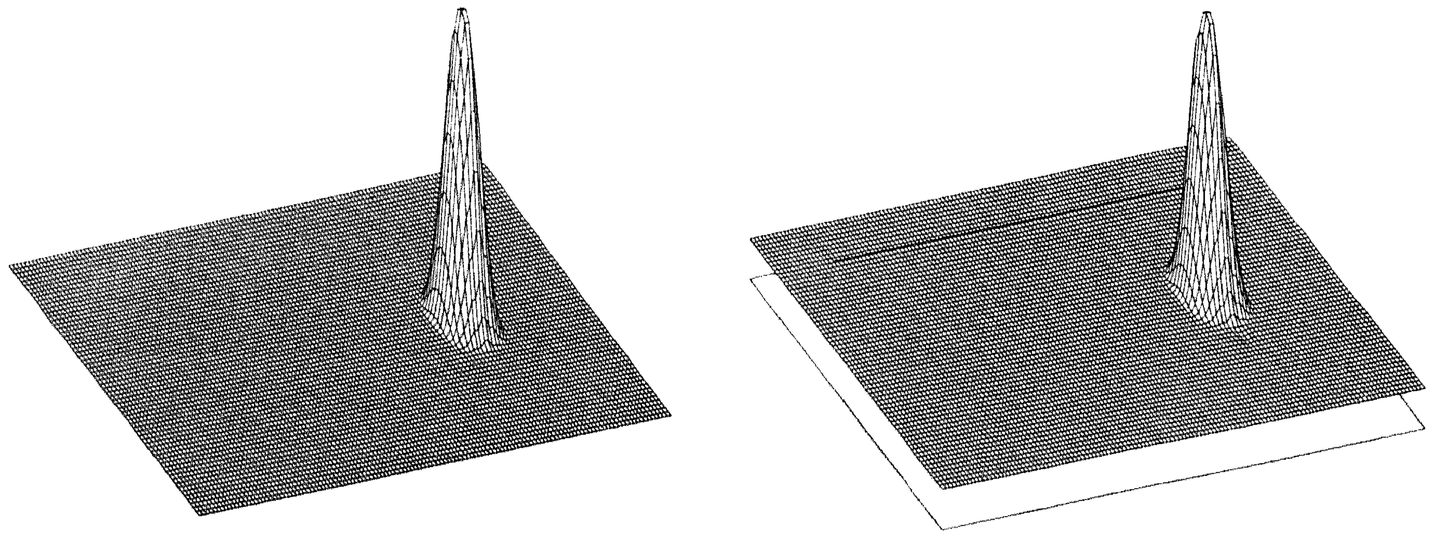

Figure 3.5: Cone tests for deformation correction.

applied on the uniform grid, but after each complete time step the values at the cells adjacent to the poles are mixed, that is, for $j = 1$ and $j = m$ we replace the $c_{i,j}^{n+1}$ by the average value

$$\frac{1}{2m} \sum_{i=1}^{2m} c_{i,j}^{n+1}. \tag{3.38}$$

We shall refer to this procedure as polar mixing. The difference with the polar caps is that we now have zonal transport inside the caps (that is, transport in the λ direction) in the first fractional step (3.26). The results are shown in Figure 3.4. Obviously, the splitting error is strongly reduced but also the diffusive effect has nearly disappeared. Although there still is a slight deformation of the background and some inaccuracy at the poles, this seems acceptable for practical purposes.

Remark. It is clear from the above figures that the modifications (3.30) and (3.31) do not give accurate results at the poles. We note, however, that without these modifications the results with either polar caps or polar mixing become much worse.

3.3.4 Deformation correction

An other way of reducing the errors near the poles follows from the observation that these errors are already present for the background concentrations. This suggests the following modification: each time the wind field is updated we compute the result of one step of the splitting scheme starting with solution equaling 1. Let us call the result of this one step $\alpha_{i,j}$. Along with this, we also take one step with a nonsplitted scheme. Starting again with solution equaling 1, we obtain the result

$$\begin{aligned} \beta_{i,j} \quad &= 1 + \frac{\Delta t}{\Delta \lambda} \left(U\left(\lambda_{i-\frac{1}{2}}, \phi_j\right) - U\left(\lambda_{i+\frac{1}{2}}, \phi_j\right) \right) \\ &\quad + \frac{\Delta t}{\Gamma_j \Delta \phi} \left(\Gamma_{j-\frac{1}{2}} V\left(\lambda_i, \phi_{j-\frac{1}{2}}\right) - \Gamma_{j+\frac{1}{2}} V\left(\lambda_i, \phi_{j+\frac{1}{2}}\right) \right). \end{aligned} \tag{3.39}$$

Next we store the factors $\delta_{i,j} = \beta_{i,j}/\alpha_{i,j}$, which will be used to correct the splitting scheme. Let the result of the uncorrected splitting step (3.26), (3.27) be denoted by $\bar{c}_{i,j}^{n+1}$. We then correct this value according to

$$c_{i,j}^{n+1} = \delta_{i,j}\, \bar{c}_{i,j}^{n+1} \tag{3.40}$$

and accept this as the new approximation at time level t_{n+1}. We note that only near the poles the $\delta_{i,j}$ will differ significantly from 1.

Figure 3.5 shows the results for the cone tests. The error near the poles now has completely disappeared, and the results are accurate. The scheme appears to be positive but not strictly monotonic. (There is a small dip trailing the cone, see also Table 3.1.)

There are two drawbacks with this approach. First, at each update of the wind field the factors $\delta_{i,j}$ have to be computed and stored. Secondly, due to the fact that this is not a modification on the fluxes, but on the $c_{i,j}$ values, the remaining scheme will no longer be strictly mass conservative. However, if there are many species to be advected, say ten or more, or if the wind field is not often updated, then the calculation of the $\delta_{i,j}$ is a relatively small task. Furthermore, in numerical tests we observed that the mass conservation is still almost valid, see for instance the results for ERR1 in Table 3.1.

It should be noted that we do not have a mathematical justification for these deformation corrections. The loss of the mass conservation property may cause difficulties in other applications. If so, the polar mixing approach of the previous section seems a good alternative.

3.3.5 Accuracy and CPU times

In Table 3.1 the errors are listed for two versions of the splitting scheme: Split-PoMx with polar mixing (Figure 3.4) and Split-DeCo with deformation corrections (Figure 3.5). Both versions are considered here with and without limiting.

The results in the table are for the cone-shaped profile (3.33) and for a cylinder of height 1, with the same base as the cone, placed on top of a background value 1. The wind field is given by (3.35), and the splitting schemes are again applied with 256 time steps for one full rotation. To enable comparison with the methods used in [SR91], the same error measures are considered and the standard first-order upwind (donor-cell) scheme has been included with 5120 time steps (maximal Courant number of about 1), which is to be compared with the MPDATA-1,1,0 scheme in [SR91]. Along with the CPU time in minutes on a SGI workstation, we also have listed the quantity CPU', which is the scaled time with respect to the donor cell scheme.

For completion we give the formulas for the error measures used in Table 3.1. They represent the errors of the minimum and maximum, the scaled L_2-error and the errors in mean value and variance, respectively. The sums and maximum-minimum

	EMIN	EMAX	ERR0	ERR1	ERR2	CPU min	CPU'
Cone tests							
donor cell	0	-0.83	0.063	0	-0.86	6.2	1
Split-PoMx (non-lim.)	-0.02	-0.15	0.011	0	-0.09	0.63	0.10
Split-DeCo (non-lim.)	-0.02	-0.12	0.009	0.001	-0.08	0.62	0.10
Split-PoMx (limited)	0	-0.17	0.011	0	-0.12	0.83	0.13
Split-DeCo (limited)	0	-0.15	0.009	0.001	-0.11	0.80	0.13
Cylinder tests							
donor cell	0	-0.30	0.067	0	-0.023	6.2	1
Split-PoMx (non-lim.)	-0.05	0.06	0.029	0	-0.004	0.63	0.10
Split-DeCo (non-lim.)	-0.03	0.07	0.028	-2×10^{-5}	-0.004	0.61	0.10
Split-PoMx (limited)	-0.06	0.02	0.028	0	-0.006	0.81	0.13
Split-DeCo (limited)	-0.02	0.02	0.028	7×10^{-5}	-0.006	0.80	0.13

Table 3.1: Results for one revolution around the sphere, 256 time steps for the splitting methods, 5120 time steps for the donor cell scheme. The values for the error measures are set to 0 if they approach single precision round-off ($< 10^{-6}$).

values are taken over $i = 1, \cdots, 2m,\ j = 1, \cdots, m$,

$$\begin{aligned}
\text{EMIN} &= \frac{\min\left(c^n_{i,j}\right) - \min\left(c^0_{i,j}\right)}{\max\left(c^0_{i,j}\right)}, \qquad \text{EMAX} = \frac{\max\left(c^n_{i,j}\right) - \max\left(c^0_{i,j}\right)}{\max\left(c^0_{i,j}\right)}, \\
\text{ERR0} &= \frac{\left(\sum \gamma_j \left(c^n_{i,j} - c^0_{i,j}\right)^2\right)^{1/2}}{\max\left(c^0_{i,j}\right)}, \qquad \text{ERR1} = \frac{\sum \gamma_j c^n_{i,j}}{\sum \gamma_j c^0_{i,j}} - 1, \\
\text{ERR2} &= \frac{\sum \gamma_j \left(c^n_{i,j}\right)^2}{\sum \gamma_j \left(c^0_{i,j}\right)^2} - 1,
\end{aligned} \tag{3.41}$$

where $\gamma_j = \Gamma_j / (2m \sum_{k=1}^{m} \Gamma_k)$. The scaling is chosen such that ERR0 will be equal to 1 if the error is 1 in all grid points and $\max(c^0_{i,j}) = 1$. (We are not completely sure whether this is the same scaling as in [SR91].)

The results for the cone show that the flux limiter only slightly affects the height of the peak. The limited schemes are positive but approximately 30% more expensive than the nonlimited counterparts.

The results for the cylinder with background clearly show that the schemes are not monotone. Even with the limited Split-DeCo scheme there still is a few percent

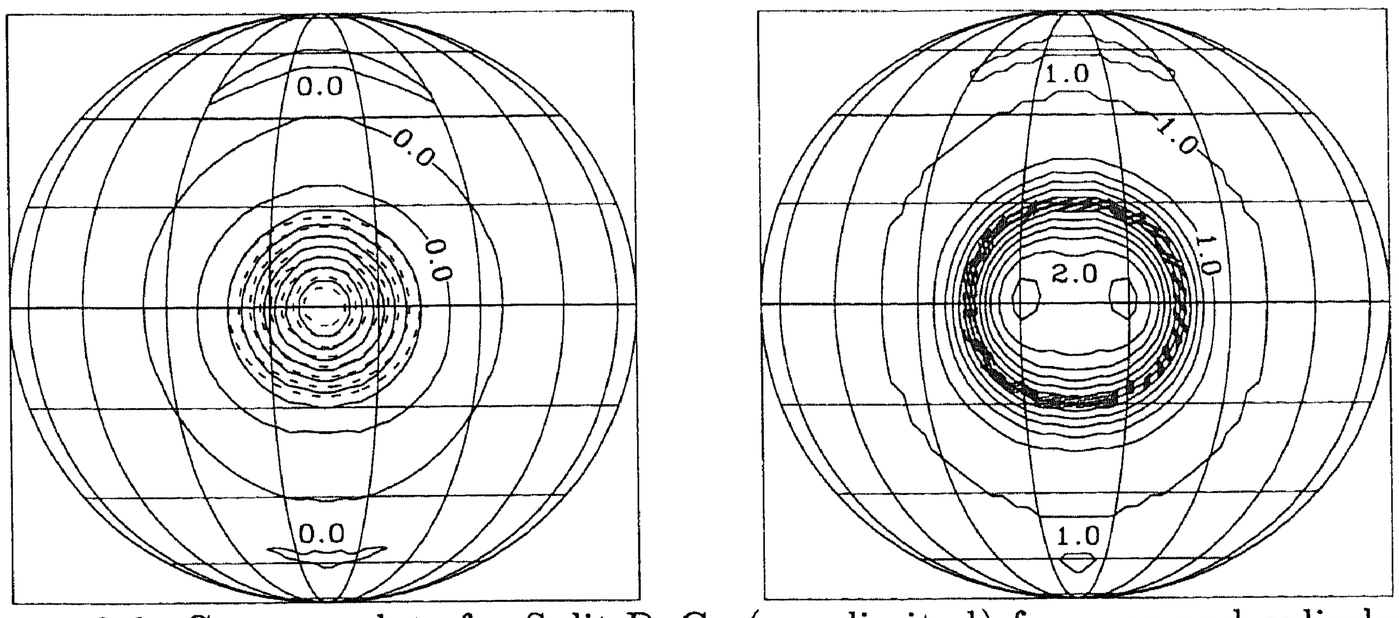

Figure 3.6: Contour plots for Split-DeCo (non-limited) for cone and cylinder.

under- and overshoot; see also the contour plots in Figure 3.7. For the Split-PoMx schemes the largest undershoots occur at the poles, in agreement with Figure 3.4.

Outside the polar regions the Split-PoMx and Split-DeCo schemes have a comparable accuracy. The Split-DeCo schemes have smaller errors at the poles but they are not strictly mass conserving.

Comparison with the semi-Lagrangian TREMBA schemes used in [SR91] shows that the splitting schemes are more accurate than the second-order method, but slightly less accurate than the nonlimited TREMBA schemes of order 4 and 6. However, with the limited versions of these schemes in [SR91] the peak height of the cone is strongly reduced and the mass errors become much worse. Also with these semi-Lagrangian schemes limiting is necessary to achieve positivity.

The main advantage of the splitting schemes seems the computational efficiency. The schemes require only a fraction of the CPU time of the donor cell scheme (with 5120 time steps), whereas the limited semi-Lagrangian schemes in [SR91] appear several times more expensive than the donor cell scheme.

For further comparison we have also included contour plots for the Split-DeCo schemes with the same figure lay-out as in [SR91]. The isolines are drawn at $0.0 - \epsilon$, $0.1 - \epsilon$, $0.2 - \epsilon$, $\cdots$ with $\epsilon = 10^{-4}$ and with solid lines for the numerical solution, dashed lines for the exact solution.

The plots for the cone reveal a small phase error, which is probably introduced at the crossing of the poles. The cylinder test in Figure 3.7 clearly shows that the scheme with limiter is not monotone, but it should be noted that the over- and undershoots are small and the cylinder shape is reasonably well maintained.

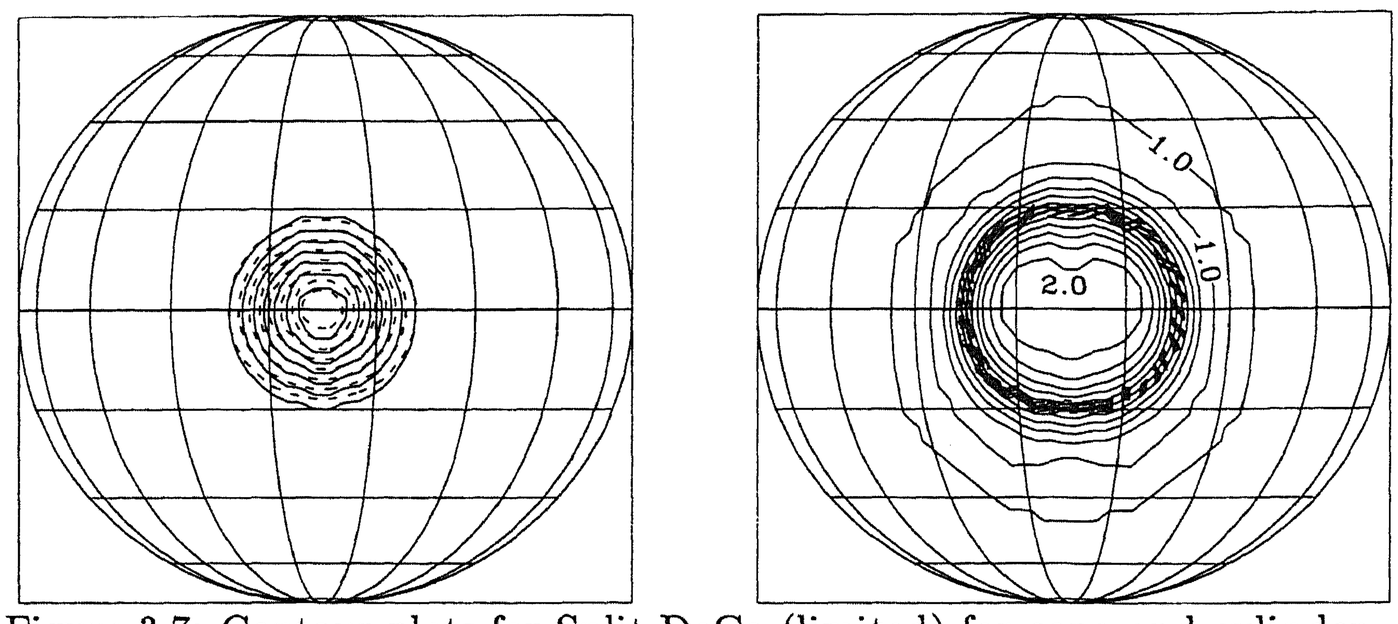

Figure 3.7: Contour plots for Split-DeCo (limited) for cone and cylinder.

In conclusion, both versions of the splitting scheme perform well on this solid-body rotation, which was designed specifically to test the suitability of numerical schemes for transport over the poles. Outside the polar regions the transport is similar to planar advection, for which we know that the schemes are efficient and accurate, see [HT94]. Due to the smaller errors at the poles the Split-DeCo schemes seem to have a slight advantage over the Split-PoMx versions. The flux limiter gives positive solutions without adding much numerical diffusion or clipping of peaks. In the near future these schemes will be tested in a more general 3D setting with nonlinear chemistry and also with reduced grids included in the considerations, similar to Rasch [Ras94].

Remark. As a further experiment we also considered the wind field (3.35) with angle

$$\beta = \pi/4, \tag{3.42}$$

and with the center of cone and cylinder shifted to $(\pi, 0)$. The profile then travels only over the North Pole but the center does not remain on the same meridian. Again, at $t = 1$ one full rotation has been completed. The numerical results were similar to those in the Figures 3.4-3.5 and Table 3.1, and therefore they are not reproduced here.

Acknowledgements The authors wish to thank Han van Dop for his comments on the presentation of the paper.

3.4 Appendix

3.4.1 Limiting for meridional transport

The limiter described in Section 3.2.1 needs a minor adjustment for the fractional step (3.27), which treats the meridional transport. Without this adjustment small negative values are created (magnitude in the range 10^{-4}–10^{-5}).

In the following we omit the index i and the asterisk in (3.27). Further we assume for the moment that $V > 0$ and that the Courant numbers are at most 1 (the large Courant numbers arise only in the λ direction). With limiting we then have [see (3.12)],

$$\begin{aligned} c_j^{n+1} = c_j \quad &+ \quad \nu_{j-\frac{1}{2}} \frac{\Gamma_{j-\frac{1}{2}}}{\Gamma_j} \left(c_{j-1} + \psi\left(\nu_{j-\frac{1}{2}}, \theta_{j-1}\right)(c_j - c_{j-1}) \right) \\ &- \quad \nu_{j+\frac{1}{2}} \frac{\Gamma_{j+\frac{1}{2}}}{\Gamma_j} \left(c_j + \frac{1}{\theta_j}\psi\left(\nu_{j+\frac{1}{2}}, \theta_j\right)(c_j - c_{j-1}) \right), \end{aligned} \tag{3.43}$$

which can be written as

$$c_j^{n+1} = Ac_j + Bc_{j-1}, \tag{3.44}$$

with

$$\begin{aligned} A \;=\; 1 \quad &+ \quad \nu_{j-\frac{1}{2}} \frac{\Gamma_{j-\frac{1}{2}}}{\Gamma_j} \psi\left(\nu_{j-\frac{1}{2}}, \theta_{j-1}\right) \\ &- \quad \nu_{j+\frac{1}{2}} \frac{\Gamma_{j+\frac{1}{2}}}{\Gamma_j} \left(1 + \frac{1}{\theta_j}\psi\left(\nu_{j+\frac{1}{2}}, \theta_j\right) \right), \end{aligned} \tag{3.45}$$

$$\begin{aligned} B \;=\; &\nu_{j-\frac{1}{2}} \frac{\Gamma_{j-\frac{1}{2}}}{\Gamma_j} \left(1 - \psi\left(\nu_{j-\frac{1}{2}}, \theta_{j-1}\right) \right) \\ &+ \quad \nu_{j+\frac{1}{2}} \frac{\Gamma_{j+\frac{1}{2}}}{\Gamma_j \theta_j} \psi\left(\nu_{j+\frac{1}{2}}, \theta_j\right). \end{aligned} \tag{3.46}$$

The requirement for positivity is $A \geq 0, \quad B \geq 0$.

Assume that $\psi(\nu, \theta) = 0$ for $\theta \leq 0$ and $0 \leq \psi(\nu, \theta) \leq 1$. This ensures that $B \geq 0$. The requirement $A \geq 0$ is satisfied if

$$\begin{aligned} \frac{1}{\theta}\psi\left(\nu_{j+\frac{1}{2}}, \theta\right) \;&\leq\; \frac{1 - \nu'_{j+\frac{1}{2}}}{\nu'_{j+\frac{1}{2}}}, \\ \nu'_{j+\frac{1}{2}} \;&=\; \frac{\Gamma_{j+\frac{1}{2}}}{\Gamma_j}\nu_{j+\frac{1}{2}}. \end{aligned} \tag{3.47}$$

This is enforced by replacing in (3.11) the expression $(1 - \nu)\theta/\nu$ by $(1 - \nu')\theta/\nu'$.

Likewise, if $V < 0$ we get the requirement

$$\begin{aligned} \frac{1}{\theta}\psi(\nu_{j-\frac{1}{2}}, \theta) \;&\leq\; \frac{1 - \nu''_{j-\frac{1}{2}}}{\nu''_{j-\frac{1}{2}}}, \\ \nu''_{j-\frac{1}{2}} \;&=\; \frac{\Gamma_{j-\frac{1}{2}}}{\Gamma_j}\nu_{j-\frac{1}{2}}, \end{aligned} \tag{3.48}$$

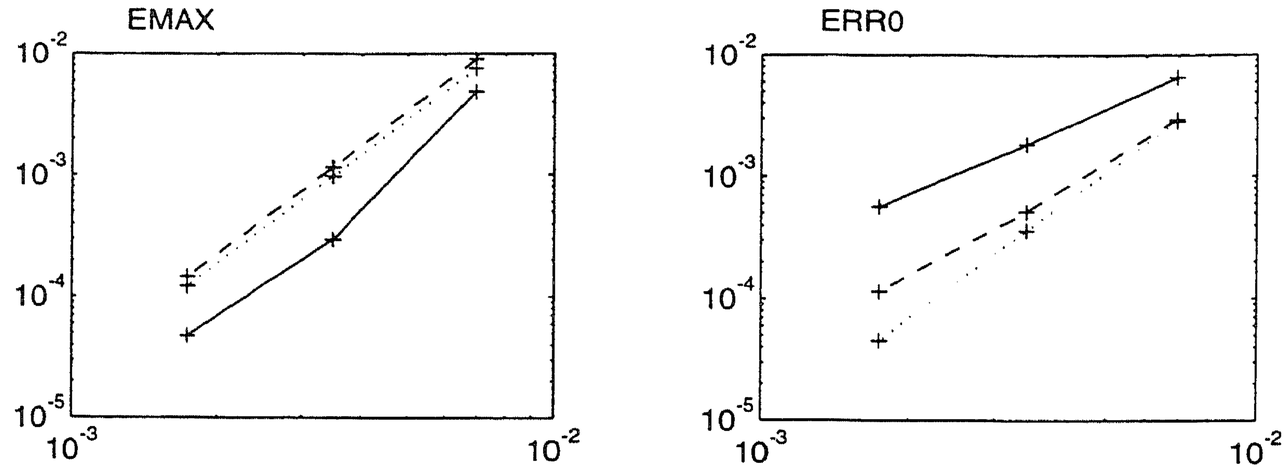

Figure 3.8: Errors for nonlimited Split-DeCo scheme versus step size 0.5 m. Results for $\beta = 0$ (dotted lines $\cdots$), for $\beta = \pi/4$ (dashed lines - - -), and for $\beta = \pi/2$ (solid lines —).

and so now we replace $(1-\nu)\theta/\nu$ in (3.11) by $(1-\nu'')\theta/\nu''$.

By a somewhat tedious calculation, it follows that the limiter will also give positivity in case we have outflow at both cell boundaries, that is, if $V_{j-\frac{1}{2}} < 0$ and $V_{j+\frac{1}{2}} > 0$, provided that both velocities are sufficiently small. This can be considered as a smoothness condition, since $V = 0$ somewhere inside the cell.

3.4.2 Convergence test

The results in Section 3.5 were given for a fixed grid size and with nonsmooth solutions. To see how the Split-DeCo scheme behaves asymptotically we present here some results for the smooth initial profile

$$c(\lambda, \phi) = \cos(\lambda - \pi/2)^4 \cos(\phi)^4 \tag{3.49}$$

on various grids.

The wind field (3.35) is considered with three values of the angle β, namely, $\pi/2$, $\pi/4$ and 0. Note that for $\beta = 0$ the wind field is parallel to the equator and constant in the λ direction, so that in this case we are in fact dealing with decoupled 1D advection problems with constant velocity, and there will be no deformation correction (all $\delta_{i,j}$ are 1). Hence the nonlimited scheme is known to be third order accurate for this case.

Below the accuracy measures EMAX and ERR0 are displayed after one rotation with the three wind fields and for three uniform $2m \times m$ grids, with $m = 36$, 72 and 144, respectively. The step size was taken as $\tau = 0.5\ m$, so that the Courant number remains the same on all grids for the $\beta = 0$ case, namely 0.5. For the other wind fields the Courant numbers are much larger, up to 45.8 for the finest grid with $\beta = \pi/2$. We note that the mass errors ERR1 were always smaller than 10^{-3} in these tests, with slow decrease upon grid refinement.

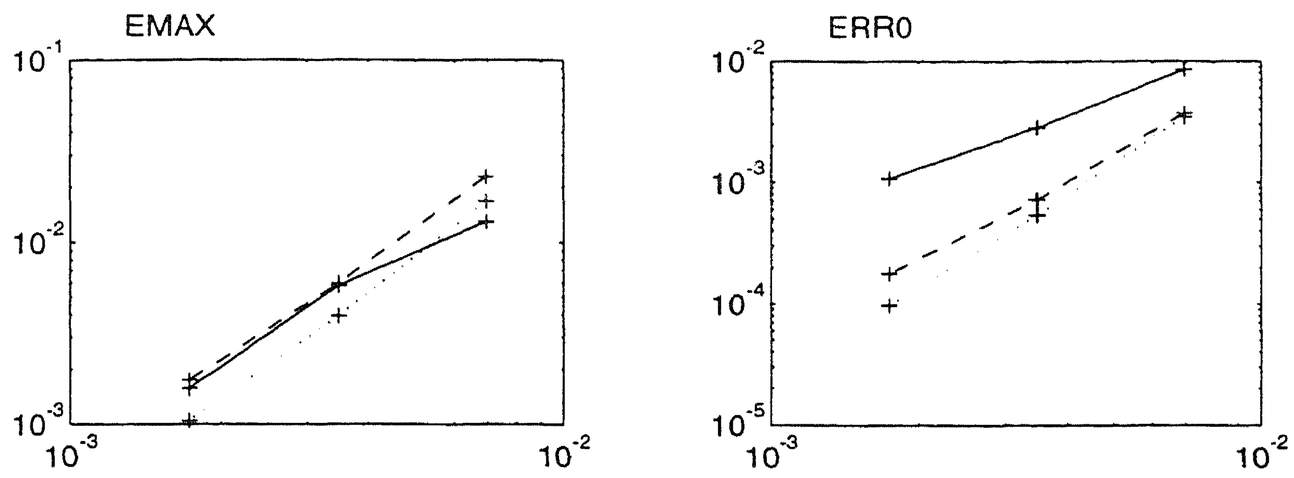

Figure 3.9: Errors for limited Split-DeCo scheme versus step size 0.5 m. Results for $\beta = 0$ (dotted lines $\cdots$), for $\beta = \pi/4$ (dashed lines - - -), and for $\beta = \pi/2$ (solid lines —).

Figure 3.8 shows the accuracy for the nonlimited Split-DeCo scheme plotted as function of the time step. As expected the convergence behaviour for the case $\beta = 0$ is best. The order of convergence for the other two wind fields is less than 3 but still satisfactory. The fact that on the coarsest grid the error EMAX is smallest for $\beta = \pi/2$ is probably due to the fact that here the peak of the profile travels over the poles, where the mesh width in λ direction strongly decreases.

Figure 3.9 gives the same results for the limited Split-DeCo scheme. With this smooth profile the limiting procedure gives some loss of accuracy, but still the results are quite accurate. Of course, limiting is only useful for profiles with sharp gradients where the nonlimited scheme will give oscillations or negative values.

Chapter 4

Coupling advection and chemical kinetics in a global atmospheric test model

Abstract

In this chapter we consider the numerical difficulties that arise when horizontal advection is coupled with chemistry on a sphere, using operator splitting. From a numerical point of view, these two processes are the most difficult parts of an atmospheric model for global studies. The advection process is solved on a uniform grid and on a so-called reduced grid, where less cells are used near the poles than at the equator. The chemical kinetics is solved with Twostep, a solver based on the two-step backward differentiation formula, with constant and variable time steps. A reproducible test to measure the accuracy and the mass conservation property is introduced. To reduce the splitting error we integrate the chemistry along characteristics.

Note: This chapter is a revision of report NM-R9508. A shortened version of this chapter has been published in *Air Pollution III, Volume 1: Air Pollution, Theory and Simulation*, Editors: H. Power and N. Moussiopoulos and C.A. Brebbia, 1995, pages 319-326.

4.1 Introduction

Global atmospheric transport models are becoming increasingly important for the analysis and description of the distribution of trace gases and particles. Examples of these models are the GISS model of Jacob et al. [JPWM87], from which the TM2 model of Heimann and Keeling [HK89] was derived, and the Moguntia model

of Zimmermann [Zim84]. Important processes in these models are advection by wind and chemical kinetics. Other processes are deposition, emission and diffusion.

In this chapter horizontal advection and chemical kinetics on a sphere are described by

$$\frac{\partial c}{\partial t} + \frac{1}{a\,\cos\phi}\left[\frac{\partial(u\,c)}{\partial\lambda} + \frac{\partial(v\,c\,\cos\phi)}{\partial\phi}\right] = F(t,\,\lambda,\,\phi,\,c), \tag{4.1}$$

where c is a concentration vector with unit particles cm^{-3}, F is the operator describing the chemical kinetics, with unit particles $\mathrm{cm}^{-3}\mathrm{s}^{-1}$, u and v are the (given) wind velocity components in respectively the λ (longitude) and ϕ (latitude) direction and a is the radius of the earth, which is taken as a perfect sphere with radius 6378 km. In global models Equation (4.1) is usually solved with operator splitting, such that advection and chemistry are handled with different numerical schemes and different time steps.

Desirable properties for the advection scheme are positivity, monotonicity and mass conservation combined with good shape preserving properties. We investigate the use of the unconditionally stable splitting scheme from [HS95] and a method of lines (MOL) scheme from [HKvLV95]. Both schemes are provided with flux limiting to minimize oscillations and negative values. These schemes compare favorably with TREMBA and MPDATA from [SR91]. On uniform grids the unconditional stability is attractive to avoid a too stringent CFL restriction emanating from the pole singularity [Wil92]. However, the pole singularity also renders a stable scheme applied with large CFL number inaccurate. To overcome this difficulty, the deformation correction proposed in [HS95] is applied. An alternative way to deal with the pole singularity is to use a so-called reduced grid [Wil92], which coarsens the grid when approaching the poles. In this paper we use both a uniform and a reduced grid. The dimension splitting scheme and the MOL scheme are applied on a reduced grid in the manner proposed in [BHV94].

The chemical solver must be computationally efficient in the 1% error range and preferably mass conservative. In this chapter, the chemical kinetics are solved with Twostep [Ver94, VS95], a solver based on the two-step backward differentiation formula, combined with Gauss-Seidel iteration to approximately solve the implicitly defined solution. For atmospheric chemistry problems Twostep compares favorably to QSSA [VS95] and VODE [VBvLS96].

The research described in this chapter deals on the one hand with the implementation of earlier derived advection schemes [BHV94, HS95] and on the other hand with the coupling of advection and chemical kinetics.

4.2 The numerical test problem

We construct a test problem for which we can calculate a reference solution with an accuracy up to round-off error. For this purpose we use an analytical wind field, so

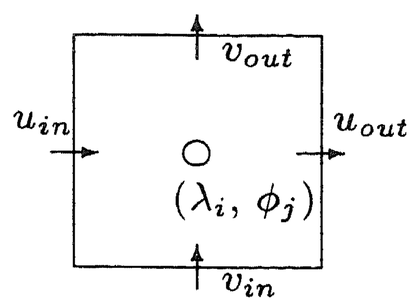

Figure 4.1: In and outgoing velocities.

that we can transform the PDE (4.1) into an ODE for all grid points by integrating along characteristics. These ODEs are solved with the state-of-the-art BDF code VODE [BBH89]. The reference solution is calculated in double precision (14 decimal digits), all other calculations in single precision (7 decimal digits).

4.2.1 The wind field

We use the wind velocities defined in [SR91]. These velocities describe a rotation over the globe, with an arbitrary angle β with the equator,

$$\begin{aligned} u(\lambda,\phi) &= (\cos\beta\cos\phi + \sin\beta\sin\phi\cos\lambda)\cdot\kappa, \\ v(\lambda,\phi) &= -\sin\beta\sin\lambda\cdot\kappa. \end{aligned} \tag{4.2}$$

We choose $\beta = 45°$ as a compromise between the most difficult numerical test ($\beta = 90°$) and the prevalent wind ($\beta = 0°$). This wind field is numerical divergent free, if the grid sizes in the λ and ϕ direction are equal, because the divergence of the solid-body wind field is proportional to

$$\begin{aligned} & -v_{out}\cos\phi_{j+\frac{1}{2}} + u_{out} + v_{in}\cos\phi_{j-\frac{1}{2}} - v_{in} \\ = \; & -\sin\beta\sin\lambda_i\cos\phi_{j+\frac{1}{2}} + \cos\beta\cos\phi_j + \sin\beta\sin\phi_j\cos\lambda_{i+\frac{1}{2}} \\ & + \sin\beta\sin\lambda_i\cos\phi_{j-\frac{1}{2}} - \left(\cos\beta\cos\phi_j + \sin\beta\sin\phi_j\cos\lambda_{i-\frac{1}{2}}\right) \\ = \; & -\sin\beta\sin\lambda_i\left(\cos\phi_{j+\frac{1}{2}} - \cos\phi_{j-\frac{1}{2}}\right) + \sin\beta\sin\phi_j\left(\cos\lambda_{i+\frac{1}{2}} - \cos\lambda_{i-\frac{1}{2}}\right) \\ = \; & \sin\beta\sin\lambda_i\left(2\sin\phi_j\sin\tfrac{\Delta\phi}{2}\right) - \sin\beta\sin\phi_j\left(2\sin\lambda_i\sin\tfrac{\Delta\lambda}{2}\right) \\ = \; & 2\sin\beta\sin\lambda_i\sin\phi_j\left(\sin\tfrac{\Delta\phi}{2} - \sin\tfrac{\Delta\lambda}{2}\right). \end{aligned} \tag{4.3}$$

See Figure 4.1 for the definition of v_{out}, u_{out}, v_{in} and u_{in}. We used the equality

$$\cos(x+\Delta x) - \cos(x-\Delta x) = -2\sin(x)\sin(\Delta x).$$

In our test, we scaled the wind such that one rotation lasts 14 days, using

$$\kappa = \frac{2\pi a}{14 \times 24 \times 3600},$$

where a is the radius of the earth and the constant $14 \times 24 \times 3600$ is just 14 days in seconds. The unit of κ is m s^{-1}. With this wind field it is possible to exactly calculate the paths of an air packet in a Lagrangian calculation. That is, for each grid point Equation (4.1) transforms into an ODE

$$\frac{dc}{dt} = F(t, \lambda^*(t), \phi^*(t), c). \tag{4.4}$$

The calculation of $\lambda^*(t)$ and $\phi^*(t)$ is given in Appendix 4.6.2. In our tests we simulate one rotation which starts at day 181 at 00.00 hours.

4.2.2 Chemical model

We use a Methane oxidation cycle which we obtained from Han The at the RIVM [The94]. This model is used in long term global studies. The model has 17 non constant species, two lump species and 46 reactions, and is stiff. The model is fully described in [VBvLS96] where it is tested as a box model. In the 2D test described in this chapter, we use realistic values for the solar zenith angle, which makes the chemistry time and space dependent. Due to the absence of day or night in the polar regions the different species grow to a maximum or decrease to a minimum. This gives sharp gradients at the transition from polar region to the normal regions (see the concentration of N_2O_5 and NO in Figure 4.2) which makes the advection process difficult. Because the total number of N atoms is conserved in the

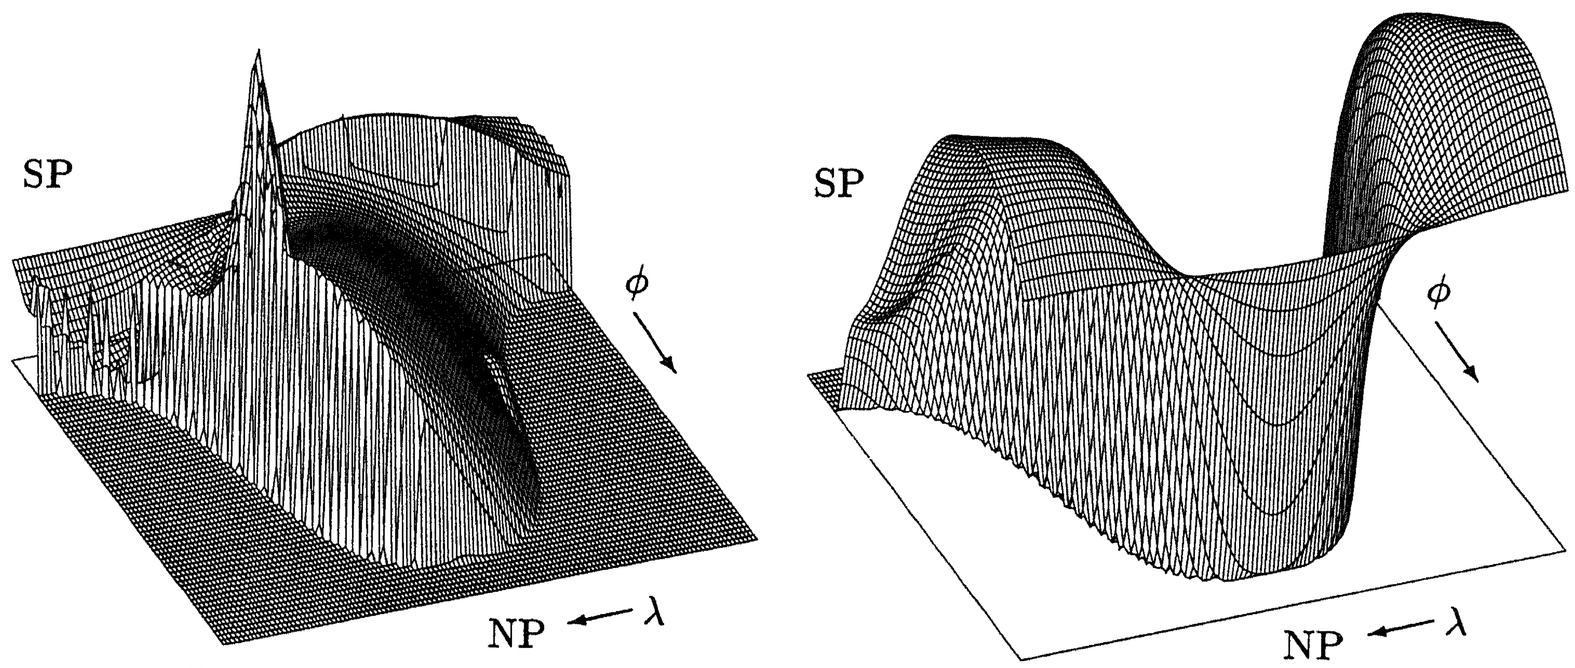

Figure 4.2: Concentration plot for $N_2 \cup_5$ (left) and NO (right)

Methane oxidation cycle, we can test the conservation properties of the numerical scheme by summing all N atoms and comparing this with the sum of the initial

value. All relevant information with respect to this chemical model can be found in `http://www.cwi.nl/ftp/edwins/CIRK_Prod_Loss.html` .

4.2.3 Initial concentration

Numerical tests for advection schemes often use a constant background concentration disturbed with, e.g., a cone or a cylinder. If all components of the initial concentration would be shaped like that, the system would be far from chemical equilibrium. Therefore our test run uses a background concentration that is in chemical equilibrium. This is achieved by starting a calculation at day 167 with the same concentration vector over the whole earth until we arrive at day 181. The initial concentrations for HNO_3 and NO are respectively 2.55×10^9 and 100 molec/cm^3, see Table 5.1 for the other species. At day 181 we then add two cones on the concentrations of HNO_3 and NO to disturb the equilibrium. We use two cones, to have always one cone in the day zone and one in the night zone,

$$\begin{aligned} c(\mathrm{HNO_3};\, t,\, \lambda,\, \phi) &= c^*(\mathrm{HNO_3};\, t,\, \lambda,\, \phi) + 4.0 \times 10^9\; c_\gamma(\lambda,\, \phi) \\ c(\mathrm{NO};\, t,\, \lambda,\, \phi) &= c^*(\mathrm{NO};\, t,\, \lambda,\, \phi) + 1.0 \times 10^9\; c_\gamma(\lambda,\, \phi) \end{aligned} \tag{4.5}$$

with

$$c_\gamma(\lambda,\, \phi) = \max\left(0,\, 1 - r^{(1)}(\lambda,\, \phi)/R,\, 1 - r^{(2)}(\lambda,\, \phi)/R\right),$$

where $R = 7\Delta\lambda$, with $\Delta\lambda$ the grid length in the λ direction, and

$$r^{(j)}(\lambda, \phi) = 2\sqrt{\cos^2\phi\, \sin^2\left[\tfrac{1}{2}\,(\lambda - \lambda_j)\right] + \sin^2\left[\tfrac{1}{2}\,(\phi - \phi_j)\right]},$$

$(\lambda_j,\, \phi_j)$ are the coordinates of the center of the day and night zone cone. These coordinates are $(\frac{\pi}{2},\, 0)$, $(1\frac{1}{2}\pi,\, 0)$.

4.3 Numerical solver

For the advection part of Equation (4.1) explicit methods are usually more efficient than implicit ones. Because of the stiffness of the chemical model, the chemical kinetics requires an implicit method. To combine these methods we consider operator splitting.

4.3.1 Computational grid

Computations are performed on two different grids, the well known uniform longitude latitude grid and a reduced grid. The specific reduced grid and data structure has been borrowed from [BHV94]. The advantages of the reduced grid are a less severe CFL restriction at the poles, so that greater time steps can be made with standard explicit schemes. A favorable side effect is the reduction of the number of grid cells in

which the chemical computations must be done. On the uniform grid in our test we use for the number of grid points in the λ and ϕ direction N_λ and N_ϕ respectively 128 and 64, resulting in a grid distance of 2.8125° in both directions. For the reduced grid we double the latitude cell width at $|\phi| = 61.875°, 75.9375°, 84.375°$, see Figure 4.3. This results in a minimum value for $\Delta\lambda \cos\phi$ which is about 5.1 times smaller than

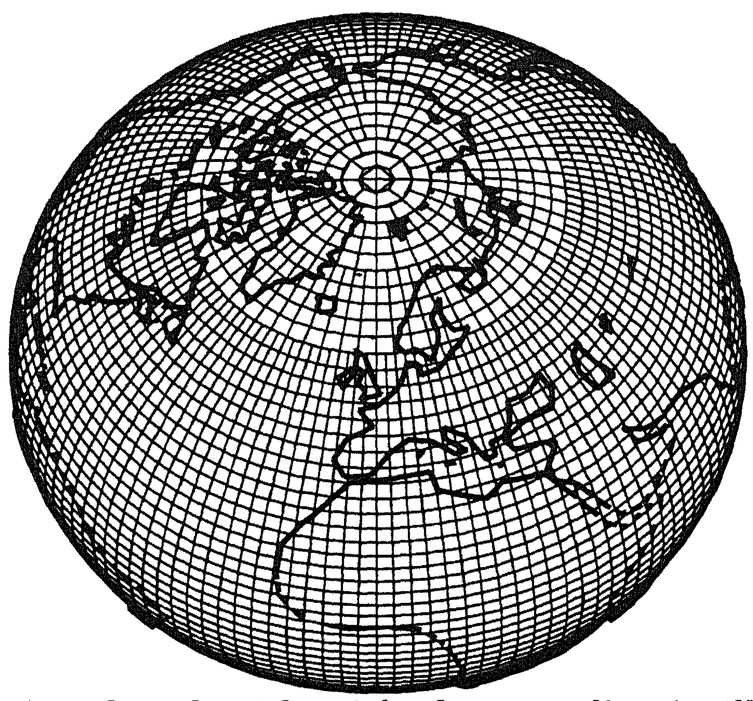

Figure 4.3: A reduced grid, with the coastline in illustration.

the value at the equator. On a uniform grid this factor would be about 40.9. The total number of grid points on a uniform grid amounts to 8192 versus 6304 on a reduced grid, which reduces the cpu time for the chemical computations with 23%.

4.3.2 Advection scheme

The advection scheme Split-DeCo introduced in [HS95] is used on the uniform longitude latitude grid and on the reduced grid. The Split-DeCo scheme is made unconditionally stable by allowing the stencil to vary with the Courant numbers. When we have a large CFL number, this scheme is still stable, but the scheme is inaccurate near the poles due to the pole singularity. To overcome this difficulty, a deformation correction is applied, see details in [HS95] The MOL scheme [HKvLV95] is used on the reduced grid only. Both schemes use the same four-point, upwind biased stencil in each direction and are provided with flux limiting to minimize the occurrence of oscillations and negative values. The MOL scheme uses a 2-stage Runge-Kutta method for the time integration and hence is approximately two times more expensive for a given step size than the Split-DeCo scheme.

4.3.3 Chemical solver

The chemical solver used is Twostep [Ver94, VS95]. Twostep is a stiff ODE-solver based on the 2nd-order implicit backward differentiation formula (BDF2). Twostep has been designed as a special purpose solver for atmospheric chemistry ODE problems and uses explicit Gauss-Seidel iteration for approximately solving the implicit relations. As in [VBvLS96] we use NO_X and O_X lumping to enhance accuracy for a low number of iterations.

To avoid many rejected steps in Twostep caused by the sharp transitions at sunrise and sunset, we wrote a subroutine which calculates sunrise and sunset, and let Twostep restart at these times. The restart is accomplished by giving Twostep sunset or sunrise as the start or end of the integration interval, if this occur. In this case we can also restart when using constant time steps. This restart might be also necessary for parametrizations in the full scale model that we have in mind, which also strongly depend on the solar zenith angle.

4.3.4 The operator splitting technique

Discretization of the advection operator in Equation (4.1) yields an ODE system,

$$\frac{\mathrm{d}c(t)}{\mathrm{d}t} = A(t, c(t)) + F(t, \lambda, \phi, c(t)).$$

We use Strang splitting. Because the space-dependency of the solar zenith angle causes a splitting error, we improve Strang splitting for this case by integrating the chemistry along characteristics. For first order splitting this correction is proposed in [HV95]. Because of the startup costs of the chemical solver, we calculate the chemistry in the middle step to have a twice as large integration interval. Thus we use the splitting

$$\begin{aligned}
&\frac{\mathrm{d}c^*(t)}{\mathrm{d}t} = A(t, c^*(t)), && t_n \le t \le t_{n+\frac{1}{2}},\\
&&& c^*(t_n) = c_n(t_n),\\
&\frac{\mathrm{d}c^{**}(t)}{\mathrm{d}t} = F\left(t, \lambda + \frac{u}{a\cos\phi}\left(t - t_{n+\frac{1}{2}}\right), \phi + \frac{v}{a}\left(t - t_{n+\frac{1}{2}}\right), c^{**}(t)\right), && t_n \le t \le t_{n+1},\\
&&& c^{**}(t_n) = c^*\left(t_{n+\frac{1}{2}}\right),\\
&\frac{\mathrm{d}c^{***}(t)}{\mathrm{d}t} = A(t, c^{***}(t)), && t_{n+\frac{1}{2}} \le t \le t_{n+1},\\
&&& c^{***}\left(t_{n+\frac{1}{2}}\right) = c^{**}(t_{n+1}),
\end{aligned} \tag{4.6}$$

giving the new approximation at time level t_{n+1}.

The meteorological input used in the full scale model we have in mind is updated every 6 hours. This implies that the advection step should be 6 hours divided by

an integer. Combined with the CFL condition in the ϕ direction for the Split-DeCo scheme on a uniform grid, this gives 224 advection steps for a full rotation. This makes the advection step 1.5 h and the chemical integration interval 3 h. To have the same integration interval for the chemistry when using the dimensional splitting scheme and the MOL scheme for advection, we divided the advection step in the MOL scheme in 4 sub advection steps of 1.5/4 h to maintain stability. This is essential for the comparison we make, because the length of the time interval $t_{n+1} - t_n$ determines the splitting error and the influence of the start up costs of the chemistry.

4.4 Numerical experiments

method	adv	shape	emin	emax	err0	err1	cpu (s)	cpu' %
Donor cell	5120	cone	0.0	-8.3E-1	6.3E-2	0.0	206	100.0
DeCo-u	256	cone	0.0	-1.5E-1	8.8E-3	-1.0E-3	20	9.8
DeCo-rg	256	cone	-2.2E-3	-1.8E-1	1.0E-2	-2.8E-3	22	10.9
MOL-rg	1450	cone	-2.4E-6	-2.9E-1	1.6E-2	0.0	150	72.9
Donor Cell	5120	cyl	0.0	-3.0E-1	6.7E-2	0.0	206	100.0
DeCo-u	256	cyl	-2.3E-2	1.9E-2	2.8E-2	7.6E-5	20	9.8
DeCo-rg	256	cyl	-1.6E-2	2.3E-2	3.0E-2	6.1E-5	22	10.9
MOL-rg	1450	cyl	-2.4E-3	1.3E-3	3.2E-2	0.0	150	72.9

Table 4.1: Results from a test with different advection schemes.

4.4.1 Numerical experiment with the advection schemes

The advection schemes are tested with a full rotation over both poles ($\beta = 90°$). By way of comparison, we also apply the donor cell scheme. The advected concentrations are a cone with no background and maximum 0.9 (cone) and a cylinder with background 1.0 and maximum 2.0 (cyl). The number of advection steps (adv) is 5120 for donor cell (as in [SR91]), 256 for Split-DeCo (as in [HS95]) 1450 and for the MOL scheme (the nearest even number as in [BHV94], because of Strang splitting we want an even number). Notice that now $\beta = 90°$ which necessitates a somewhat larger number of steps than necessary for $\beta = 45°$. The results are given in Table 4.1, where emin and emax are respectively the error of the minimum and the maximum value, err0 is the scaled L_2-error, and err1 is a measure for the mass balance. The

errors are the same as in [HS95, SR91],

$$\text{emin} = \frac{\min\left(c^n_{i,j}\right) - \min\left(c_{i,j}(t_n)\right)}{\max\left(c_{i,j}(t_n)\right)}, \qquad \text{emax} = \frac{\max\left(c^n_{i,j}\right) - \max\left(c_{i,j}(t_n)\right)}{\max\left(c_{i,j}(t_n)\right)},$$

$$\text{err0} = \frac{\left(\sum \gamma_j (c^n_{i,j} - c_{i,j}(t_n))^2\right)^{1/2}}{\max\left(c_{i,j}(t_n)\right)}, \qquad \text{err1} = \frac{\sum \gamma_j c^n_{i,j}}{\sum \gamma_j c_{i,j}\left(t_n\right)} - 1,$$

where

$$\gamma_j = \cos\left(\phi_j\right) / \left(N_\lambda \sum_{k=1}^{N_\phi} \cos(\phi_k) \right).$$

The scaling is chosen such that err0 will be equal to 1 if the error is 1 in all grid points and $\max(c_{i,j}(t_n)) = 1$. The sums and max/min values are taken over $i = 1, \cdots, N_\lambda,\ j = 1, \cdots, N_\phi$. In all formulas $c_{i,j}(t_n)$ denotes the reference solution, and $c^n_{i,j}$ the computed solution at time t_n. cpu' is the relative cpu time compared to the donor cell scheme. The results are for double precision to overcome some difficulties caused by underflow due to numerical diffusion.

The err0 error of the Split-DeCo and the MOL scheme are of the same order of magnitude, and comparable with the errors of the scheme in [Ras94], but this scheme seems to be more expensive in terms of computational effort. The MOL scheme uses more cpu time and is more diffusive, but is closer to physics as it uses no splitting. The Split-DeCo schemes are not strictly mass conservative, but the err1 error is acceptable. The schemes on a reduced grid are not positive, but this can easily be repaired in cases where positivity is very important, see Appendix 4.6.1. The Split-DeCo scheme gives the best results, even on a uniform grid.

4.4.2 Numerical experiment with Twostep

To examine the time step strategy in Twostep we have first integrated along characteristics, where errors are only caused by the chemical solver. We use 2 Gauss-Seidel iterations and in the tests with variable time steps between τ_{min} and τ_{max}, the relative tolerances (rtol) are 10^{-1}, 10^{-2} and 10^{-3}. The absolute tolerance is always 1.0. A lower bound of $\tau_{min} = 30$ s prevents Twostep from taking time steps which are too small. Increasing this lower bound is strongly dissuaded, because the BDF2 solution will then sometimes give unphysical negative concentrations. The results are given in Table 4.2, where L2(err0) is the L_2-error over the components over err0, the global L_2-error. $\overline{\text{err1}}$ is the mean over the components over err1 and is a measure for the mass balance.

Although variable time steps give good results, we have also performed runs with constant time steps, because this will probably be more attractive when working on a vector machine. Then, 120 s gives an optimum between cpu time and accuracy for constant time steps, although the BDF2 solution then sometimes gives negative

τ_{min}, τ_{max} (s)	rtol	L2(err0)	$\overline{\text{err1}}$	cpu (h)
1, 900	1E-1	4.8E-3	9.0E-3	1.87
30, 900	1E-1	4.8E-3	9.0E-3	1.72
1, 900	1E-2	2.4E-3	4.5E-3	3.85
30, 900	1E-2	2.4E-3	4.5E-3	2.13
1, 900	1E-3	2.6E-4	4.6E-4	10.03
240, 240	–	1.7E-2	2.9E-2	1.61
120, 120	–	5.7E-3	9.6E-3	3.13
60, 60	–	7.4E-4	1.4E-3	6.22

Table 4.2: Results from a test along characteristics.

solutions. Using a constant time step of 30 s in this test costs about 12 hours. Results for rtol $= 10^{-1}$ and $\tau_{min} = 30$ s gives for the lowest cpu time the desired accuracy of approximately 1%.

4.4.3 Advection and chemical kinetics

method	τ_{min}, τ_{max} (s)	rtol	L2(err0)	L2(emax)	$\overline{err1}$	cpu (s)
DeCo-u	30, 900	1E-1	8.7E-3	2.5E-2	8.5E-3	1.05
DeCo-rg	30, 900	1E-1	8.9E-3	2.5E-2	8.4E-3	0.88
MOL-rg	30, 900	1E-1	9.6E-3	4.6E-2	8.3E-3	1.17
DeCo-u	1, 900	1E-3	7.5E-3	3.1E-2	2.2E-4	6.21
DeCo-rg	1, 900	1E-3	7.8E-3	3.1E-2	2.2E-4	5.59
MOL-rg	1, 900	1E-3	8.6E-3	5.3E-2	1.5E-4	5.82

Table 4.3: Advection with chemical kinetics on a uniform and a reduced grid.

In Table 4.3 we show the results of coupling the Split-DeCo scheme and the MOL scheme with Twostep using variable time steps. We use the same errors as in Table 4.2, extended with an error measure representing smearing, L2(emax), the L_2-error over the components over emax, the error in the maximum.

In all cases the mass balance error is mainly determined by the chemistry solver, whereas the smearing is mainly determined by the advection solver. The global L_2-error over all components is caused by the operator splitting, the advection solver and the chemistry solver, although in the case with the most accurate chemistry, this

term is negligible. The Split-DeCo scheme on a reduced grid gives for the lowest cpu time the desired accuracy. The MOL scheme on a reduced grid costs more cpu time and is more diffusive.

In Table 4.4 results are given for a test without a restart at sunset and sunrise, and results for a calculation without integrating chemistry along the characteristics. The advection is carried out with the Split-DeCo scheme on a uniform grid. From this table it is clear that a restart at sunrise and sunset is redundant. On the other hand, integration along characteristics reduces the global L_2-error from approximately 1.4% to 0.8%, without extra cpu time.

Figure 4.4 and 4.5 are plots for the sum of N-atoms, after a full rotation on a reduced grid. There is no splitting error visible at the poles with the Split-DeCo scheme, and it is clear that the MOL approach is somewhat more diffusive, in accordance with the tests presented in [HS95, HKvLV95].

restart	along characteristics	τ_{min}, τ_{max} (s)	rtol	L2(err0)	L2(emax)	$\overline{err1}$	cpu (h)
yes	yes	1, 900	1E-3	7.5E-3	2.8E-2	2.4E-4	6.21
no	yes	1, 900	1E-3	7.5E-3	3.1E-2	2.7E-4	6.23
yes	no	1, 900	1E-3	1.3E-2	2.9E-2	1.5E-3	6.32
no	no	1, 900	1E-3	1.3E-2	2.9E-2	1.5E-3	6.46
yes	yes	30, 900	1E-1	8.7E-3	2.5E-2	8.5E-3	1.05
no	yes	30, 900	1E-1	8.7E-3	2.5E-2	8.6E-3	1.09
yes	no	30, 900	1E-1	1.4E-2	2.4E-2	9.0E-3	1.02
no	no	30, 900	1E-1	1.4E-2	2.4E-2	9.1E-3	1.05
yes	yes	120, 120	–	9.6E-3	2.4E-2	1.1E-2	2.31
no	yes	120, 120	–	9.7E-3	2.3E-2	1.1E-2	2.28
yes	no	120, 120	–	1.5E-2	2.2E-2	1.4E-2	2.26
no	no	120, 120	–	1.5E-2	2.2E-2	1.4E-2	2.27

Table 4.4: Effect of restart and integration along characteristics.

4.5 Results and conclusions

The tests described in this chapter show that the reduced grid is a good approach to reduce cpu time, while maintaining accuracy. However, they reveal that the Split-DeCo scheme also performs very satisfactorily on the uniform grid due to the unconditional stability in the longitudal direction to overcome the pole singularity, as suggested in

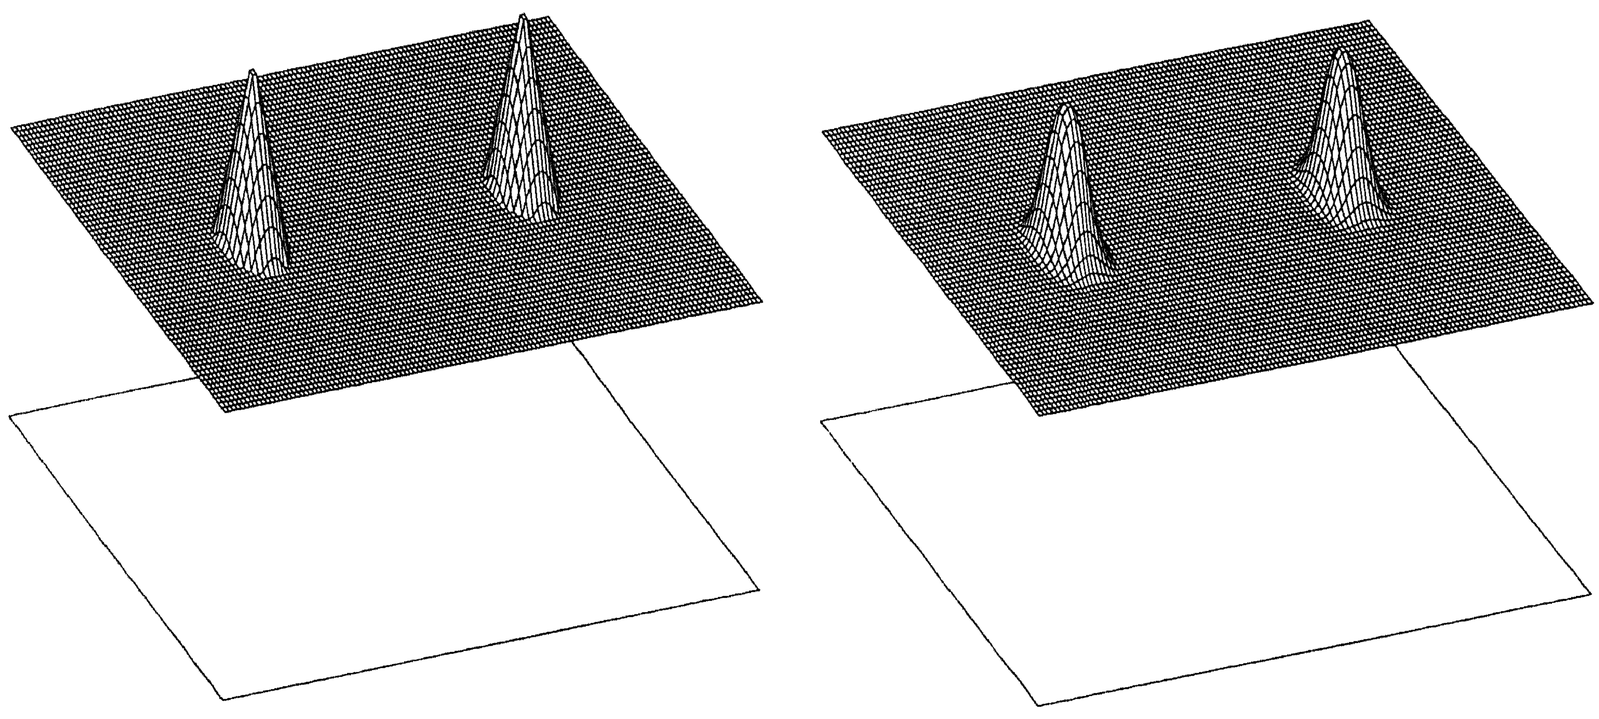

Figure 4.4: Reference solution (left) and redgridDeCo (right)

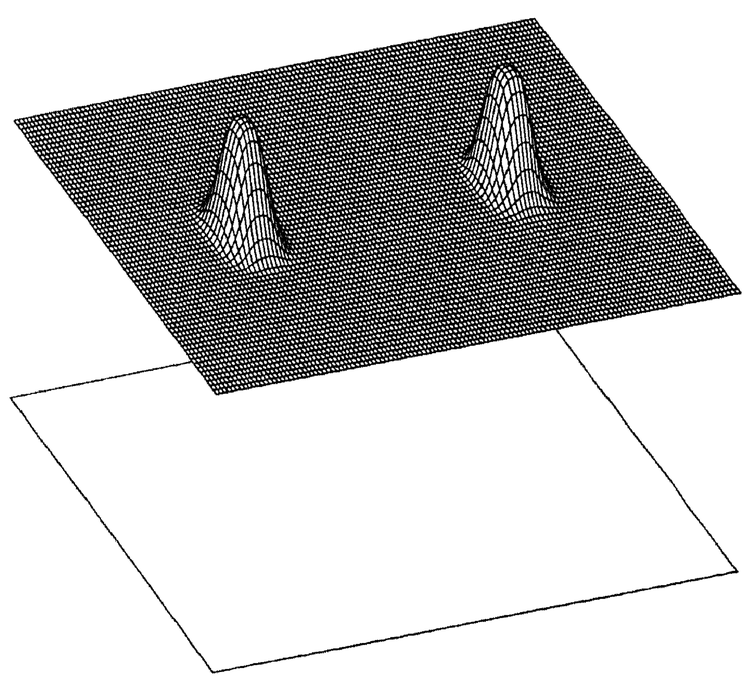

Figure 4.5: Results for the MOL approach.

[HS95]. On the other hand, grid reduction is necessary for the MOL scheme for coping with singularities at the poles.

On the scalar machine used here, variable time steps in the chemistry solver give the best results. We will test whether this is also valid on vector and parallel machines in the near future. All combinations of schemes in this chapter give the desired accuracy of approximately 1%, see Tables 4.3 and 4.4, except the smearing for the MOL scheme. The cpu time for the Split-DeCo scheme is dominated by the chemistry. In the MOL approach the cpu time is caused for roughly 63% by the chemistry and for 37% by the advection. The MOL approach costs roughly 33% more cpu time than the Split-DeCo approach and is more diffusive. Coupled with chemistry, the Split-DeCo scheme performs well on both a uniform and a reduced grid. The splitting error caused by the time splitting of the advection and chemistry is about 0.6% for the split step size of 3 h. This error can be reduced by integrating chemistry along characteristics.

4.6 Appendix

4.6.1 Implementation of the reduced grid

Comments on the implementation of the reduced grid are from [BHV94].

The reduced grid is implemented in a domain decomposition fashion where uniform blocks of different latitude cell widths are connected. At the interfaces the required concentrations are obtained by interpolation. To preserve mass conservation we impose as boundary conditions for the coarser blocks the fluxes computed on the boundary of the neighboring finer domains. If we assume one grid coarsening at $\phi_{J+\frac{1}{2}}$, and consequently also at $\phi_{N_\phi - J+\frac{1}{2}}$, we would get the following implementation:

- First handle the equatorial area, i.e., compute the fluxes for $i = 1, \cdots, N_\lambda$ and $j = J+1, \cdots, N_\phi - J$. If values c_{iJ} or $c_{iN_\phi - J+1}$ are needed they are obtained by linear interpolation from the available values

$$\begin{aligned} c_{ij} &= \frac{1}{4}\left(3c_{\frac{i}{2}j} + c_{\frac{i}{2}+1j}\right), \quad i \text{ even} \\ c_{ij} &= \frac{1}{4}\left(3c_{\frac{i}{2}+1j} + c_{\frac{i}{2}j}\right), \quad i \text{ odd} \end{aligned} \tag{4.7}$$

- Next, compute the fluxes at the boundary for the South-Pole area

$$\begin{aligned} \mathrm{fP}_{iJ+\frac{1}{2}} &= \frac{1}{2}\left(\mathrm{fP}_{2i-1J+\frac{1}{2}} + \mathrm{fP}_{2iJ+\frac{1}{2}}\right) \quad i = 1, \cdots, N_\lambda/2 \\ \mathrm{fP}_{i\frac{1}{2}} &= 0 \end{aligned}$$

- Finally, compute the the fluxes at the boundary for the North-Pole area

$$\begin{aligned} \mathrm{fP}_{iN_\phi - J+\frac{1}{2}} &= \frac{1}{2}\left(\mathrm{fP}_{2i-1N_\phi - J+\frac{1}{2}} + \mathrm{fP}_{2iN_\phi - J+\frac{1}{2}}\right) \quad i = 1, \cdots, N_\lambda/2 \\ \mathrm{fP}_{iN_\phi+\frac{1}{2}} &= 0 \end{aligned}$$

The implementation for more domains is analogous.

The spatial discretization scheme implemented as above can result in negative values. This can be avoided by taking constant interpolation at the boundary of coarse to finer blocks, instead of linear interpolation. That means that Equation (4.7) is replaced by

$$\begin{aligned} c_{ij} &= c_{\frac{1}{2}j}, & i \text{ even} \\ c_{ij} &= c_{\frac{1}{2}+1j,}, & i \text{ odd} \end{aligned} \tag{4.8}$$

This leads to an order reduction at the boundary of coarse to finer blocks. After the several tests we have done so far, it is not clear whether positivity is more important than the influence of the order reduction.

4.6.2 Exact solution of the advection

In this section we calculate the position of a vector (λ_0, ϕ_0) at time τ, the time divided by the time of a full rotation, after being advected by the velocities given in Equation (4.2) from [SR91],

$$\begin{aligned} u &= (\cos\beta\,\cos\phi + \sin\beta\,\sin\phi\,\cos\lambda)\cdot\kappa \\ v &= -\sin\beta\,\sin\lambda\cdot\kappa. \end{aligned}$$

These velocities describe a rotation over the earth with an angle β with the equator. The exact solution is a translation in a shifted coordinate system, which we denote with $\overline{x}, \overline{y}, \overline{z}$ (Cartesian) and $\overline{\lambda}, \overline{\phi}$ (polar). For simplicity we use coordinates in which the radius of the earth is scaled to 1. We first translate the coordinates λ_0 and ϕ_0 into $\overline{\lambda_0}$ and $\overline{\phi_0}$,

$$\begin{aligned} x_0 &= \cos\phi_0\,\cos\lambda_0 \\ y_0 &= \cos\phi_0\,\sin\lambda_0 \\ z_0 &= \sin\phi_0 \\ \overline{x_0} &= z_0\,\sin\beta + x_0\,\cos\beta \\ \overline{y_0} &= y_0 \\ \overline{z_0} &= z_0\,\cos\beta - x_0\,\sin\beta \\ \overline{\phi_0} &= \arcsin(\overline{z_0}) \\ \overline{\lambda_0} &= \arctan2\pi(\overline{y_0},\,\overline{x_0}), \end{aligned} \tag{4.9}$$

with $\arctan2\pi(y,\,x)$ is defined as

$$\begin{cases} \tan\left(\arctan2\pi\,(y,\,x)\right) = y\,/\,x \\ 0 \le \arctan2\pi\,(y,\,x) < 2\pi \end{cases} \tag{4.10}$$

In this system we can calculate the rotation in one time step due to the wind.

$$\overline{\lambda_\tau} = \overline{\lambda_0} + \tau\,2\,\pi \quad \text{and} \quad \overline{\phi_\tau} = \overline{\phi_0}. \tag{4.11}$$

Finally we translate the rotated coordinates back in the original coordinate system.

$$\begin{aligned}
\overline{x_\tau} &= \cos\overline{\phi_\tau}\,\cos\overline{\lambda_\tau} \\
\overline{y_\tau} &= \cos\overline{\phi_\tau}\,\sin\overline{\lambda_\tau} \\
\overline{z_\tau} &= \sin\overline{\phi_\tau} \\
x_\tau &= -\overline{z_\tau}\,\sin\beta + \overline{x_\tau}\,\cos\beta \\
y_\tau &= \overline{y_\tau} \\
z_\tau &= \overline{z_\tau}\,\cos\beta + \overline{x_\tau}\,\sin\beta \\
\phi_\tau &= \arcsin(z_\tau) \\
\lambda_\tau &= \arctan 2\pi(y_\tau,\, x_\tau).
\end{aligned} \tag{4.12}$$

Chapter 5

A numerical study for global atmospheric transport-chemistry problems

Abstract

Atmospheric air quality modeling relies in part on numerical simulation. Required numerical simulations are often hampered by lack of computer capacity and computational speed. This problem is most severe in the field of global modeling where transport and exchange of trace constituents are studied in the whole of the global troposphere/stratosphere. Studies in this field easily lead to computations with millions of unknowns over long time spans. Use of fast, tailored numerical algorithms implemented on the most advanced computer systems is then needed for making real progress. This chapter is devoted to a performance study of algorithms developed for solving global atmospheric transport-chemistry problems. The algorithms are applied within two different operator splittings and are vectorized and parallelized on a Cray C90. For the performance study we use a constructed, three-space dimensional model problem, containing advection, vertical turbulent diffusion and chemical reactions. This numerical model problem is presented as a benchmark on which other algorithms and implementations can be tested. The benchmark problem together with all the algorithm software is available through World Wide Web. In the context of operator splitting, particular attention is paid to the issue of solving chemistry and vertical turbulent diffusion coupled or uncoupled.

Note: This chapter is a revision of report MAS-R9702. The research described in this chapter was supported by Cray Research, Inc., under Grant 96.03, via the Stichting Nationale Computerfaciliteiten (National Computing Facilities Foundation, NCF).

5.1 Introduction

Air quality models are used to enhance the understanding of the chemical composition of the atmosphere, in particular with regard to the relation between emissions and the resulting distributions of primary and secondary polluting species. These models are mass balances which appear in the form of systems of time-dependent, three-space dimensional, partial differential equations (PDEs). These systems split additively into various subsystems describing advective transport, turbulent/diffusive transport, chemical transformations, emissions and depositions. Research into these models is becoming more and more important as it now seems evident that human activities leading to atmospheric air pollution can entail the danger of a long lasting global environmental change. The recent review papers [CSP+96, PBC+95] discuss the current state and future directions in air quality modeling.

Atmospheric air quality modeling relies in part on numerical simulation as the PDEs used are nonlinear and cannot be solved by analytical means. However, numerical simulations are often severely hampered by lack of computer capacity and computational speed. This problem is most severe in the field of global modeling where transport and exchange of trace constituents are studied in the whole of the global troposphere/stratosphere. Studies in this field easily lead to computations with millions of unknowns over long time spans. In such cases use of the most advanced computer systems is a prerequisite for making real progress in modeling research.

Bearing the practical problem of computational speed in mind, this chapter is devoted to a performance study of algorithms tailored for solving global atmospheric transport-chemistry problems. The algorithms are applied and compared within two different operator splittings and are vectorized and parallelized on a Cray C90. Performance results are presented with respect to (grid) vectorization and parallelization based on autotasking combined with a division into subdomains of the globe.

For the performance study we have used a constructed, three-space dimensional model problem containing advection, vertical turbulent diffusion and a photochemical scheme consisting of 45 reactions between 17 species. This particular model problem has allowed us to examine the numerical accuracy/efficiency of the algorithms and the operator splittings. In the context of operator splitting, particular attention is paid to the issue of solving chemistry and vertical turbulent diffusion coupled or uncoupled. The model is restricted in the sense that it does not simulate a real atmosphere. However, for numerical testing purposes it is valuable, certainly with regard to the chemistry which is based on a set of photochemical reactions from practice. The model is therefore presented as a benchmark on which other schemes and implementations can be tested as well, since we consider benchmarking important for algorithm and code development.

This chapter is organized as follows. In Section 5.2 we discuss the spherical mass balance equation used in our investigation. Section 5.3 is devoted to the numerical algorithms and the two operator splittings. The benchmark problem is discussed in Section 5.4. Section 5.5 presents the results on vectorization, showing Mflop rates

and CPU times, and discusses accuracy and efficiency. Section 5.6 is devoted to the issue of parallelization. The final Section 5.7 summarizes conclusions and remarks.

5.2 The transport-chemistry model

At the heart of atmospheric air quality models lie mass balances in the form of systems of advection-diffusion-reaction equations

$$\frac{\partial c_l}{\partial t} + \vec{\nabla} \cdot (\vec{u} c_l) = \vec{\nabla} \cdot \left(\rho D \vec{\nabla} \left(\frac{c_l}{\rho} \right) \right) + R_l(c), \qquad l = 1, \ldots, N_c. \tag{5.1}$$

The unknown c denotes a vector of species concentrations, say of length N_c, and ρ is the density of the air. The velocity wind field vector $\vec{u}$ and the diffusion matrix D are supposed to be known (off-line model). The wind field is generated separately by circulation models or retrieved from data bases and the diffusion is usually determined by atmospheric turbulence models. While advection and diffusion determine the transport of species, the reaction term R represents the atmospheric chemical reactions, depositions and emissions. Note that the reaction term couples all equations in the model. All processes are time and space dependent, but this dependence is suppressed in our notation.

The vector function $R(c)$ has the special form

$$R(c) = P(c) - L(c)\, c, \tag{5.2}$$

where $P(c)$ is the vector of production terms and $L(c)c$ the vector of loss terms with $L(c)$ a diagonal matrix. For many species, the reciprocal of their entry in $L(c)$ is a good approximation of the physical time constant or characteristic reaction time. In virtually all applications, the range of reaction times is huge, ranging from milliseconds or shorter (e.g. OH radical) to years (e.g. CH_4). This means that atmospheric chemistry gives rise to stiffness. A second important feature is photochemistry, giving rise to rapid changes in concentration values at sunrise and sunset. We thus have to face constantly moving areas of rapid solution change. In applications the number of species varies. Current global models use about 20 species, but in [CSP+96, PBC+95] it is pointed out that as many as 40 to 100 species are necessary for an adequate analysis of perturbations to atmospheric chemistry on a global and regional scale. These review papers also point out that grid resolutions of $1° \times 1°$ or better in the horizontal and at least 20 vertical grid layers are needed to achieve this. Altogether this means that global and regional air quality modeling leads to a huge computational task. Even with high-performance computers at hand, computer capacity often dictates the grid resolution.

As far as transport is concerned, in this chapter we restrict ourselves to horizontal transport by wind and vertical transport by turbulent diffusion. This means that vertical advection and horizontal diffusion is neglected. In reality these two processes

are of less importance and adding them at a later stage will not lead to large numerical difficulties. A greater restriction is that we here also neglect orography, which means that the earth is taken to be a real sphere. Nor do we use real meteorological data and hence are not confronted with massive I/O operations.

We write $c = c(t, \lambda, \phi, r)$ where $\lambda \in [0, 2\pi]$ denotes longitude, $\phi \in [-\frac{\pi}{2}, +\frac{\pi}{2}]$ latitude and $r > 0$ the radial distance. With the above restrictions in mind, we put $\vec{u} = (u, v)$ and write the divergence in spherical coordinates,

$$\vec{\nabla} \cdot (\vec{u} c_l) = \frac{1}{a \cos\phi} \left[\frac{\partial (u\, c_l)}{\partial \lambda} + \frac{\partial (v\, c_l \cos\phi)}{\partial \phi} \right], \quad l = 1, \ldots, N_c. \tag{5.3}$$

Here a is the radius of the earth (6378 km) and we should mention that we have approximated the radial distance factor $1/(r + a)$ by $1/a$. For the applications this is allowed, since the atmospheric layer (troposphere/stratosphere) in which models are used is extremely thin compared with the radius of the earth. The diffusion term becomes

$$\vec{\nabla} \cdot \left(\rho\, D\, \vec{\nabla} \left(\frac{c_l}{\rho} \right) \right) = \frac{\partial}{\partial r} \left(\rho\, K\, \frac{\partial}{\partial r} \left(\frac{c_l}{\rho} \right) \right), \quad l = 1, \ldots, N_c, \tag{5.4}$$

where K is a scalar, vertical diffusion coefficient. We suppose for r the interval $0 \le r \le r_H$ with r_H a height uniform over the globe. As boundary conditions we use the no-flux conditions

$$\rho\, K\, \frac{\partial}{\partial r} \left(\frac{c}{\rho} \right) = 0, \quad r = 0, r_H. \tag{5.5}$$

Boundary conditions in the horizontal directions are not needed of course. The initial condition at the initial time $t = t_0$ is imposed by prescribing all species concentrations.

5.3 Operator splitting and discretization

As outlined above, mathematical air quality models split additively into advective transport, turbulent/diffusive transport, chemical transformations, emissions and depositions. Virtually all computer implementations employ this in the numerical solution process by using the concept of operator (time) splitting. Splitting is popular because the submodels are different in nature and are therefore easier to solve apart than when kept together in the numerical treatment. However, splitting also has a number of disadvantages. First, it introduces an additional error, called the splitting error. This error exists even if submodels would be solved exactly, while in practice this error is hard to relate with the discretization errors. Second, splitting can introduce strong initial transients for the fast reacting species in the chemistry computation and does this within each split time step. Due to the nonlinearity, it may be necessary to resolve these transients accurately in order to avoid instabilities and inaccuracies which cannot be corrected at later times. Resolving the transients

accurately is costly, however. We thus see that it is of practical interest to apply operator splitting in a way that both these disadvantages are somehow reduced when necessary.

In this chapter we compare two splitting procedures. The first one is standard, while in the second we avoid splitting the chemistry and vertical diffusion, applying the same numerical technique as in [VBH96]. A related technique can be found in [KW93]. Keeping the chemistry and vertical diffusion coupled in the numerical solution process is natural, since both these processes contain small and large time scales (vertical turbulent diffusion transports species as fast as many react with one another). Comparative results in [GM91] for a number of 1D diffusion-chemistry models clearly indicate that uncoupling chemistry and vertical turbulent diffusion can lead to substantial errors. However, a similar investigation [Kes95] has not confirmed their findings. It therefore is of interest to address the question of decoupling chemistry and vertical diffusion.

In the remainder of this section we will first describe the two splitting procedures and after that we will present the algorithms for the submodels.

5.3.1 Operator splitting

We write Equation (5.1) as

$$\frac{\partial}{\partial t}c = F_0(c) + F_1(c) + F_2(c), \tag{5.6}$$

where

$$\begin{aligned} F_0(c) &= R(c), \\ F_1(c) &= \frac{\partial}{\partial r}\left(\rho K \frac{\partial}{\partial r}\left(\frac{c}{\rho}\right)\right), \\ F_2(c) &= -\frac{1}{a\cos\phi}\left[\frac{\partial(u\,c)}{\partial\lambda} + \frac{\partial(v\,c\cos\phi)}{\partial\phi}\right]. \end{aligned}$$

Let Δt_{split} denote the split step size and c_s the approximation to the exact concentration vector c, resulting from the splitting, at time $t = t_s$ where $t_s = s\,\Delta t_{\text{split}}$ for $s = 0, 1, \ldots$. Starting from c_s, the approximation c_{s+1} at the next time point t_{s+1} is then computed with the standard splitting procedure by solving

$$\frac{\partial}{\partial t}c^{(1)}(t) = F_2(c^{(1)}(t)) \qquad (t_s \le t \le t_{s+\frac{1}{2}}), \qquad c^{(1)}(t_s) = c_s, \tag{5.7a}$$

$$\frac{\partial}{\partial t}c^{(2)}(t) = F_1(c^{(2)}(t)) \qquad (t_s \le t \le t_{s+\frac{1}{2}}), \qquad c^{(2)}(t_s) = c^{(1)}(t_{s+\frac{1}{2}}), \tag{5.7b}$$

$$\frac{\partial}{\partial t}c^{(3)}(t) = F_0(c^{(3)}(t)) \qquad (t_s \le t \le t_{s+1}), \qquad c^{(3)}(t_s) = c^{(2)}(t_{s+\frac{1}{2}}), \tag{5.7c}$$

$$\frac{\partial}{\partial t}c^{(4)}(t) = F_1(c^{(4)}(t)) \qquad (t_{s+\frac{1}{2}} \le t \le t_{s+1}), \quad c^{(4)}(t_{s+\frac{1}{2}}) = c^{(3)}(t_{s+1}), \tag{5.7d}$$

$$\frac{\partial}{\partial t}c^{(5)}(t) = F_2(c^{(5)}(t)) \quad (t_{s+\frac{1}{2}} \le t \le t_{s+1}), \quad c^{(5)}(t_{s+\frac{1}{2}}) = c^{(4)}(t_{s+1}), \quad (5.7e)$$

and $c_{s+1} \equiv c^{(5)}(t_{s+1})$. By means of (5.7), the original problem thus has been splitted into five subproblems, of which three are different and can be treated with different numerical techniques. Note that the sequence of computations, in the order advection/vertical diffusion/chemistry/vertical diffusion/advection, is symmetrical around the time point halfway (Strang splitting [Str68]). This way of splitting is standard in the sense that all processes are treated uncoupled in a sequential manner. In the remainder, this standard splitting procedure will be called the type I splitting procedure. Problem (5.7c) will be referred to as the chemistry problem, problem (5.7a) and (5.7e) as the advection problem, and problem (5.7b) and (5.7d) as the diffusion problem.

The second splitting procedure, where the chemistry and vertical diffusion remain coupled, reads

$$\frac{\partial}{\partial t}c^{(1)}(t) = F_2(c^{(1)}(t)) \ (t_s \le t \le t_{s+\frac{1}{2}}), \ c^{(1)}(t_s) = c_s, \quad (5.8a)$$

$$\frac{\partial}{\partial t}c^{(2)}(t) = F_0(c^{(2)}(t)) + F_1(c^{(2)}(t)) \ (t_s \le t \le t_{s+1}), \ c^{(2)}(t_s) = c^{(1)}(t_{s+\frac{1}{2}}), \quad (5.8b)$$

$$\frac{\partial}{\partial t}c^{(3)}(t) = F_2(c^{(3)}(t)) \ (t_{s+\frac{1}{2}} \le t \le t_{s+1}), \ c^{(3)}(t_{s+\frac{1}{2}}) = c^{(2)}(t_{s+1}), \quad (5.8c)$$

and $c_{s+1} \equiv c^{(3)}(t_{s+1})$. This splitting procedure is nonstandard and will henceforth be called the type II splitting procedure. Problem (5.8b) will be referred to as the chemistry-diffusion problem.

In both splittings, the step size Δt_{split} must be chosen somehow in relation with the discretization errors made in the substeps. In general this choice is hard to make and requires experimental insight. Logical is to take the minimal value $\Delta t_{\text{split}} = 2\,\Delta t_{\text{adv}}$, where Δt_{adv} denotes the step size for the advection computation. We have done this in all tests presented in this chapter

5.3.2 The advection computation

For both splittings the advection computation goes the same. Because there is only horizontal advection, it suffices to describe this computation for the 2D equation

$$\frac{\partial c}{\partial t} = -\frac{1}{a\cos\phi}\left[\frac{\partial(u\,c)}{\partial\lambda} + \frac{\partial(v\,c\,\cos\phi)}{\partial\phi}\right], \quad (5.9)$$

where $c = c(t,\lambda,\phi)$ is now a scalar rather than a vector. A so-called reduced space grid [Wil92] is used, see Figure 5.1 for an example. Grid reduction means that at a small number of latitudes near the poles the grid size in the longitudinal direction is doubled. Without this doubling, a realistic grid on the globe would become very

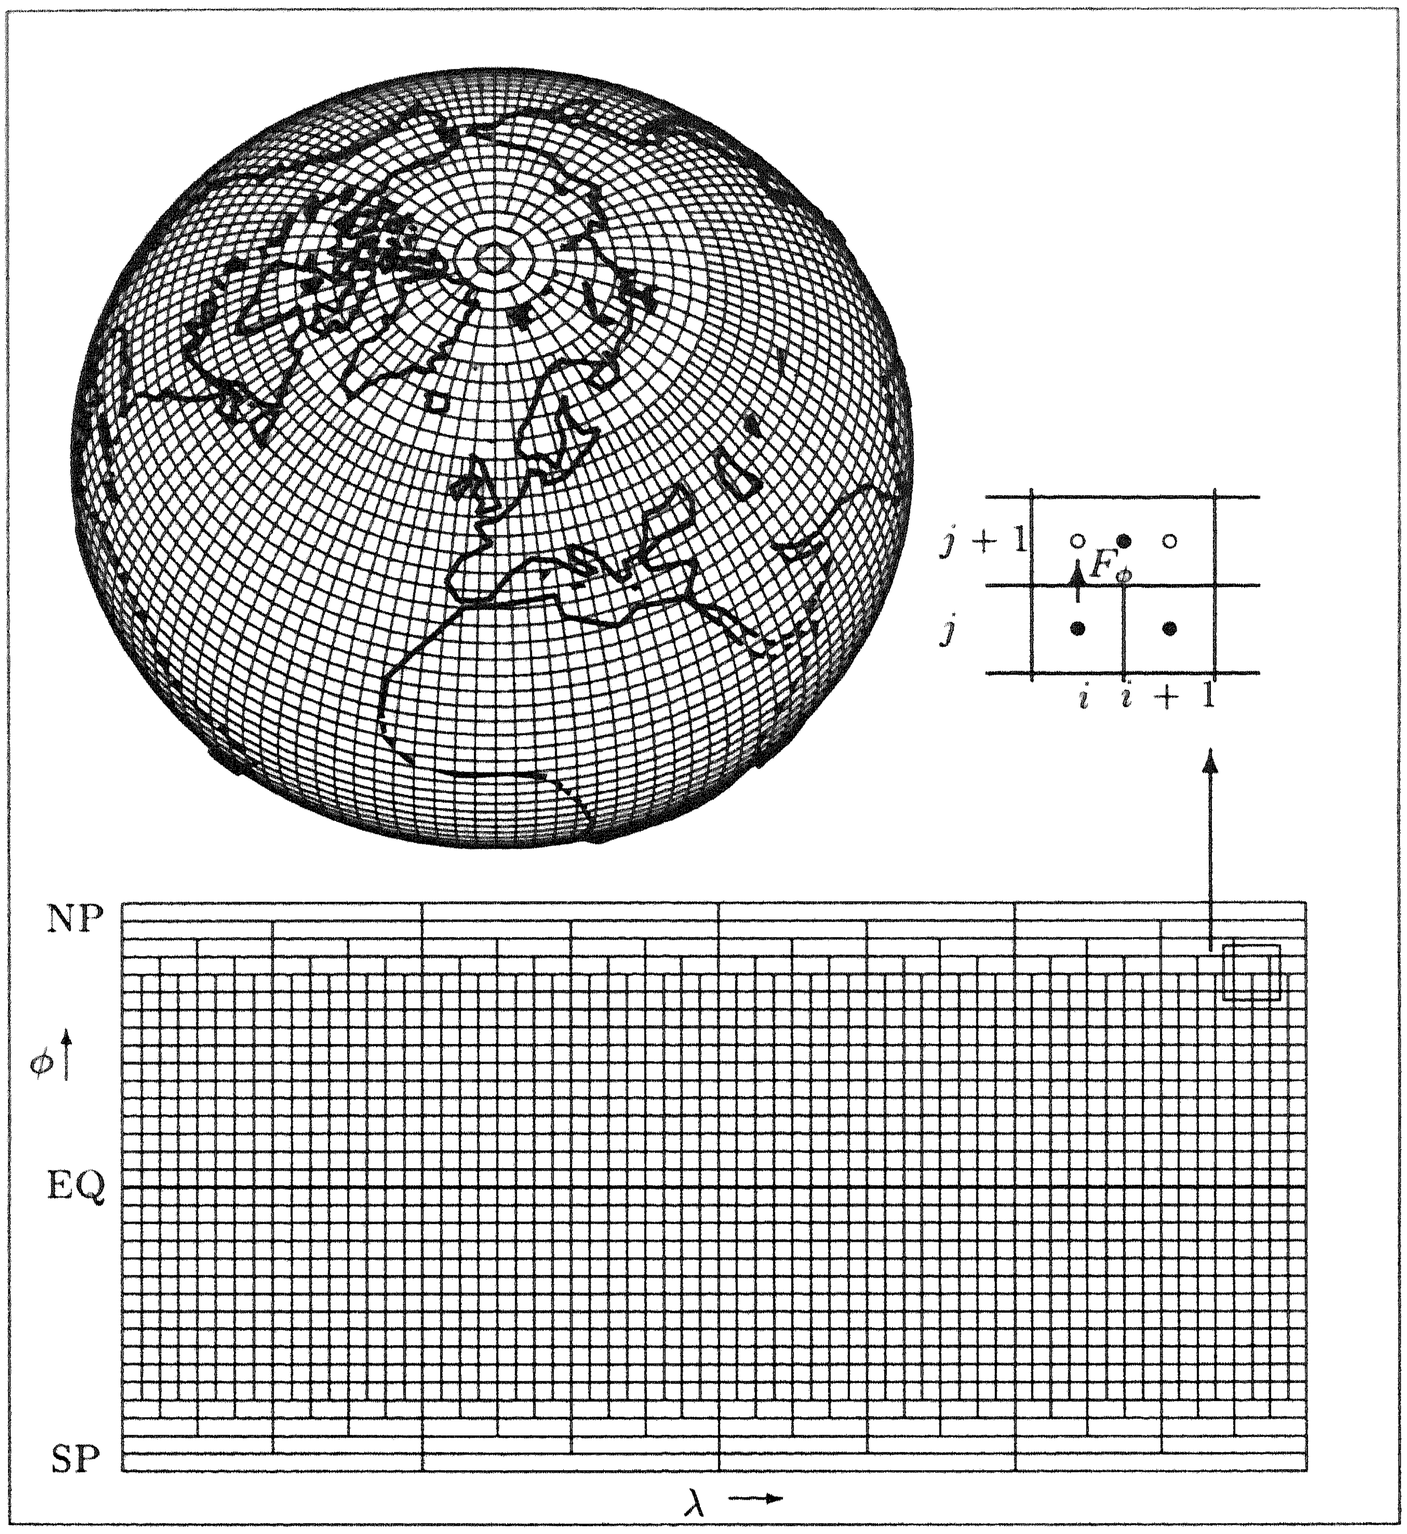

Figure 5.1: A reduced grid on the globe (left, top), virtual concentrations ∘ and real concentrations • (right, top), and a 64×32 reduced grid in longitude/latitude coordinates (bottom).

fine in the longitudinal direction near the poles, which imposes a severe stability restriction for explicit advection schemes.

The spherical advection operator is cell-centered discretized by a mass-conservative, flux-limited, third-order upwind scheme. Let $c_{i,j}(t)$ denote the resulting approximation at a cell-center (λ_i, ϕ_j) with cell size $\Delta\lambda(j) \times \Delta\phi$ ($\Delta\lambda$ depends on the latitude index j due to the grid reduction). The semi-discrete counterpart of (5.9) then reads

$$\frac{\mathrm{d}}{\mathrm{d}t} c_{i,j} = -\frac{1}{a \cos\phi_j} \left[\frac{\mathrm{fL}_{i+\frac{1}{2},j} - \mathrm{fL}_{i-\frac{1}{2},j}}{\Delta\lambda(j)} + \frac{\mathrm{fP}_{i,j+\frac{1}{2}} - \mathrm{fP}_{i,j-\frac{1}{2}}}{\Delta\phi} \right], \tag{5.10}$$

with the fluxes on the cell boundaries

$$\mathrm{fL}_{i+\frac{1}{2},j} = \max\left(0, u_{i+\frac{1}{2},j}\right) \mathrm{fL}^{+}_{i+\frac{1}{2},j} + \min\left(0, u_{i+\frac{1}{2},j}\right) \mathrm{fL}^{-}_{i+\frac{1}{2},j}, \tag{5.11a}$$

$$\mathrm{fP}_{i,j+\frac{1}{2}} = \left[\max\left(0, v_{i,j+\frac{1}{2}}\right) \mathrm{fP}^{+}_{i,j+\frac{1}{2}} + \min\left(0, v_{i,j+\frac{1}{2}}\right) \mathrm{fP}^{-}_{i,j+\frac{1}{2}}\right] \cos\phi_{j+\frac{1}{2}}. \tag{5.11b}$$

These fluxes are defined by

$$\mathrm{fL}^{+}_{i+\frac{1}{2},j} = c_{i,j} + \Psi\left(\theta_{i+\frac{1}{2},j}\right) \cdot \left(c_{i+1,j} - c_{i,j}\right), \tag{5.12a}$$

$$\mathrm{fL}^{-}_{i+\frac{1}{2},j} = c_{i+1,j} + \Psi\left(\theta^{-1}_{i+\frac{3}{2},j}\right) \cdot \left(c_{i,j} - c_{i+1,j}\right), \tag{5.12b}$$

$$\mathrm{fP}^{+}_{i,j+\frac{1}{2}} = c_{i,j} + \Psi\left(\theta_{i,j+\frac{1}{2}}\right) \cdot \left(c_{i,j+1} - c_{i,j}\right), \tag{5.12c}$$

$$\mathrm{fP}^{-}_{i,j+\frac{1}{2}} = c_{i,j+1} + \Psi\left(\theta^{-1}_{i,j+\frac{3}{2}}\right) \cdot \left(c_{i,j} - c_{i,j+1}\right), \tag{5.12d}$$

with

$$\theta_{i+\frac{1}{2},j} = \frac{c_{i,j} - c_{i-1,j}}{c_{i+1,j} - c_{i,j}}, \quad \theta_{i,j+\frac{1}{2}} = \frac{c_{i,j} - c_{i,j-1}}{c_{i,j+1} - c_{i,j}}, \tag{5.13}$$

and Ψ is the flux-limiter function [Kor93]

$$\Psi(\theta) = \max\left(0, \min\left(1, \theta, \tfrac{1}{3} + \tfrac{\theta}{6}\right)\right). \tag{5.14}$$

We used this advection approximation earlier in [BHV94, HKvLV95, Spe95]. In [HKvLV95] the particular type of flux-limiting is discussed. Flux-limiting serves to maintain positivity. If we switch off the flux-limiting, i.e. put $\Psi(\theta) = \frac{1}{3} + \frac{\theta}{6}$, the underlying third-order upwind scheme is recovered.

Our reduced grid approach was developed in [BHV94] where we refer to for a detailed discussion of the technicalities encountered. Only minor differences exist between the current implementation and the one developed in [BHV94]. For example, where grid reduction takes place, piecewise constant interpolation is used for concentration values whenever needed versus linear interpolation in [BHV94]. Piecewise constant interpolation was found to work equally well.

At this stage of development, (5.10) is still time-continuous. For the time integration an explicit Runge-Kutta method is used, viz. the second-order, two-stage explicit trapezoidal rule. Let us represent (5.10) by the system

$$\frac{\mathrm{d}}{\mathrm{d}t}\vec{c} = \vec{f}(\vec{c}), \tag{5.15}$$

where $\vec{c} = (c_{i,j})$ stands for the entire grid function for the concentrations and $\vec{f} = (f_{i,j})$ for the entire grid function formed by the right-hand sides of (5.10), here denoted by $f_{i,j}$. The complete advection scheme is then given by

$$\vec{w} = \vec{c}_m + \Delta t_{\text{adv}} \vec{f}(\vec{c}_m), \tag{5.16a}$$

$$\vec{c}_{m+1} = \vec{c}_m + \frac{1}{2}\Delta t_{\text{adv}} \left(\vec{f}(\vec{c}_m) + \vec{f}(\vec{w}) \right), \tag{5.16b}$$

where $\vec{c}_m$ approximates $\vec{c}(t)$ at a time point $t = t_m$ and $\Delta t_{\text{adv}} = t_{m+1} - t_m$ is the advection step size. The time points t_m lie in a split interval $[t_s, t_{s+1}]$ as introduced previously. Usually, Δt_{adv} is equal to Δt_{split} divided by an even number. In all tests presented in this paper we have used $\Delta t_{\text{adv}} = \Delta t_{\text{split}}/2$.

The use of this second-order two-stage Runge-Kutta method leads to two evaluations of the advection operator within each advection step. Noteworthy is that we have turned a linear advection problem into a nonlinear semi-discrete system (5.15) by applying flux-limiting. This leads to additional costs for the advection computation. However, in [HKvLV95] it is shown that the combination, formed by the flux-limited third-order discretization and this Runge-Kutta method, combines good stability with good positivity properties (see in particular Sections 3.3, 5.3 of [HKvLV95]). The results presented there indicate that the combination is stable and positive, as long as Δt_{adv} is adjusted to satisfy the CFL restriction

$$\max_{i,j} (\nu_{\lambda_{i,j}} + \nu_{\phi_{i,j}}) \leq \tfrac{2}{3}, \tag{5.17}$$

where

$$\nu_{\lambda_{i,j}} = \frac{\Delta t_{\text{adv}}}{\Delta\lambda(j)} \frac{\left| u\left(\lambda_{i+\frac{1}{2}}, \phi_j\right)\right|}{a\cos\phi_j}, \quad \nu_{\phi_{i,j}} = \frac{\Delta t_{\text{adv}}}{\Delta\phi} \frac{\left| v\left(\lambda_i, \phi_{j+\frac{1}{2}}\right)\right| \cos\phi_{j+\frac{1}{2}}}{a\cos\phi_j}. \tag{5.18}$$

5.3.3 The chemistry-diffusion computation used in type II splitting

Both the chemistry and diffusion computation used in type I splitting can be seen as simplified cases of the chemistry-diffusion computation used in type II splitting. We therefore first present the computation for the chemistry-diffusion problem

$$\frac{\partial}{\partial t}c = \frac{\partial}{\partial r}\left(\rho K \frac{\partial}{\partial r}\left(\frac{c}{\rho}\right)\right) + R(c), \quad 0 \leq r \leq r_H, \tag{5.19}$$

subjected to the boundary conditions (5.5), and recall that this computation has to be carried out over all split intervals $[t_s, t_{s+1}]$ at any of the cell centers (λ_i, ϕ_j) of the horizontal grid.

The usual numerical approach for such stiff problems leads us to implicit ODE solvers employing a form of Newton iteration for solving the implicit relations. However, despite the 1D nature, Newton iteration is still quite expensive here, as c is vector valued (N_c species). In [VB96, VBH96] an alternative technique has been proposed which for tropospheric models is significantly more efficient. This technique is based on a Gauss-Seidel iteration, which treats the chemistry in a scalarly implicit way and the diffusion implicitly. As a result, at most linear tridiagonal systems of algebraic equations need to be solved, rather than the much larger banded linear systems arising in the Newton process. We have adopted this technique for the current coupled chemistry-diffusion computation. For reasons of self-containedness, we here repeat the description of [VBH96], Section 2.

It suffices to consider the 1D case so that $c = c(t, r) \in \mathrm{R}^{N_c}$. Let us introduce the mixing ratio vector $\chi(t, r) = c(t, r)/\rho(t, r)$. The diffusion term is discretized on the nonuniform cell-centered grid

$$\Omega_V = \{r_k : r_1 = \tfrac{1}{2}\Delta r_1,\ r_k = r_{k-1} + \tfrac{1}{2}(\Delta r_{k-1} + \Delta r_k),\ 2 \le k \le N_r\} \tag{5.20}$$

in the following way,

$$\frac{\partial}{\partial r}\left(\rho K \frac{\partial}{\partial r}\left(\frac{c}{\rho}\right)\right) \approx \frac{2}{r_{k+1} - r_{k-1}}\left((\rho K)_k^+ \frac{\chi_{k+1} - \chi_k}{r_{k+1} - r_k} - (\rho K)_k^- \frac{\chi_k - \chi_{k-1}}{r_k - r_{k-1}}\right), \tag{5.21}$$

for $1 \le k \le N_r$, with $\chi_k(t)$ denoting the semi-discrete approximation to $\chi(t, r_k)$ and

$$(\rho K)_k^{\pm} = (\rho K)\left(t, (r_k + r_{k\pm1})/2\right), \quad r_0 = r_1 - \Delta r_1, \quad r_{N_r+1} = r_{N_r} + \Delta r_{N_r}.$$

Note that ρK is evaluated halfway between the cell centers, rather than at the cell boundaries, to obtain a discretization with order of consistency at least one on a nonuniform grid. The boundary conditions are incorporated by putting $(\rho K)_k^- = 0$ for $k = 1$ and $(\rho K)_k^+ = 0$ for $k = N_r$.

For the remainder it is important to recall that the diffusion operator introduces no coupling between different species. The species are coupled only through the chemistry system

$$R(c_k) = P(c_k) - L(c_k)\, c_k.$$

Let $c_k^{(j)}$ denote the j-th (component) species of c_k and introduce the following species vectors on Ω_V,

$$\mathbf{c}^{(j)} = [c_1^{(j)}, \ldots, c_{N_r}^{(j)}]^T, \quad \mathbf{P}^{(j)}(\mathbf{c}) = [P^{(j)}(c_1), \ldots, P^{(j)}(c_{N_r})]^T, \quad j = 1, \ldots, N_c. \tag{5.22}$$

The vector $\mathbf{c}$ is supposed to contain all vectors $\mathbf{c}^{(j)}$. Assume a similar definition for the diagonal matrices

$$\mathbf{L}^{(j)}(\mathbf{c}) = \mathrm{diag}(L^{(j)}(c_1), \ldots, L^{(j)}(c_{N_r})), \quad j = 1, \ldots, N_c. \tag{5.23}$$

Then we may represent the semi-discrete ODE system resulting from the spatial discretization, by

$$\frac{\mathrm{d}}{\mathrm{d}t}\,\mathbf{c}^{(j)} = \mathbf{A}\mathbf{c}^{(j)} + \mathbf{P}^{(j)}(\mathbf{c}) - \mathbf{L}^{(j)}(\mathbf{c})\mathbf{c}^{(j)}, \quad j = 1, \ldots, N_c, \tag{5.24}$$

where $\mathbf{A}$ is the tridiagonal diffusion matrix of order N_r.

This ODE system has to be integrated in time over each split interval $[t_s, t_{s+1}]$ introduced previously. At the beginning of each interval $[t_s, t_{s+1}]$, the integration is started with the well-known, first-order, one-step implicit Euler rule. Thereafter the second-order, two-step implicit backward differentiation (BDF) formula is used. This combination yields second-order accurate time stepping which for atmospheric transport applications is sufficient in view of the modest accuracy requirement. Generally, a relative accuracy larger than 1% is superfluous. In our tests described later we have integrated with a variable step size within each split interval. The step size is hereby governed by a standard local error control mechanism similar as used in [Ver94, VBH96]. Step sizes are taken variable so as to allow smaller step sizes in the initial part of the split intervals, where we may encounter rapid transient solution components connected with short living chemical species. However, we always impose a minimum for the step size (specified later), since we wish to avoid adjustment to the smallest time constants of radicals.

The two-step BDF formula is defined as follows. Let $\Delta t_{\mathrm{cvd}} = t_{n+1} - t_n$ denote the step size, assuming that $t_n \in [t_s, t_{s+1}]$ is a time point for the BDF formula. Then we have

$$\mathbf{c}_{n+1}^{(j)} = \mathbf{C}^{(j)} + \gamma \Delta t_{\mathrm{cvd}} \left(\mathbf{A}\mathbf{c}_{n+1}^{(j)} + \mathbf{P}^{(j)}(\mathbf{c}_{n+1}) - \mathbf{L}^{(j)}(\mathbf{c}_{n+1})\mathbf{c}_{n+1}^{(j)} \right), \quad j = 1, \ldots, N_c, \tag{5.25}$$

where the scalar γ and the vector $\mathbf{C}^{(j)}$ are BDF quantities defined by

$$\gamma = \frac{1+q}{1+2q}, \quad \mathbf{C}^{(j)} = \frac{1}{1+2q}\left((1+q)^2 \mathbf{c}_n^{(j)} - q^2 \mathbf{c}_{n-1}^{(j)} \right), \quad q = \frac{t_{n+1} - t_n}{t_n - t_{n-1}}. \tag{5.26}$$

If we put $\gamma = 1$ and $\mathbf{C}^{(j)} = \mathbf{c}_n^j$, the implicit Euler rule is recovered.

The BDF formula (5.25) is implicit and thus we need an iterative technique to approximately solve the $N_c \times N_r$ dimensional systems of nonlinear algebraic equations. As mentioned above, for this purpose we use a special Gauss-Seidel iteration which we describe next. For convenience of notation, we suppress the time index $n+1$ and put $\tau = \Delta t_{\mathrm{cvd}}$, so that (5.25) now reads

$$\mathbf{c}^{(j)} = \mathbf{C}^{(j)} + \gamma\tau\mathbf{A}\mathbf{c}^{(j)} + \gamma\tau\mathbf{P}^{(j)}(\mathbf{c}) - \gamma\tau\mathbf{L}^{(j)}(\mathbf{c})\mathbf{c}^{(j)}, \quad j = 1, \ldots, N_c. \tag{5.27}$$

This system is equivalent to

$$\mathbf{c}^{(j)} = \left(\mathbf{I} - \gamma\tau\mathbf{A} + \gamma\tau\mathbf{L}^{(j)}(\mathbf{c}) \right)^{-1} \left(\mathbf{C}^{(j)} + \gamma\tau\mathbf{P}^{(j)}(\mathbf{c}) \right), \quad j = 1, \ldots, N_c, \tag{5.28}$$

since the inverse of the tridiagonal matrix $\mathbf{I} - \gamma\tau\mathbf{A} + \gamma\tau\mathbf{L}^{(j)}(\mathbf{c})$ always exists. The Gauss-Seidel iteration for approximating $\mathbf{c}^{(j)}, 1 \leq j \leq N_c$, is carried out on Equation (5.28) and consists of the following calculations. Let $\mathbf{c}_{[i]}$ denote the i-th iterate for $\mathbf{c}$. Then, at integration step n, we have

1. Initial estimation: $i = 0$, $\mathbf{c}_{[i]} = \max(0, \mathbf{c}^n + q(\mathbf{c}^n - \mathbf{c}^{n-1}))$.
   2. Compute, in the order $j = 1, \ldots, N_c$:
      2*a*. $\mathbf{L}^{(j)}(\mathbf{c}_{[i]}), \mathbf{P}^{(j)}(\mathbf{c}_{[i]})$.
      2*b*. LU-decompose $\mathbf{I} - \gamma\tau\mathbf{A} + \gamma\tau\mathbf{L}^{(j)}(\mathbf{c}_{[i]})$. (5.29)
      2*c*. Backsolve $\mathbf{c}^{(j)}_{[i+1]}$.
      2*d*. Update $\mathbf{c}_{[i]} = (\mathbf{c}^{(1)}_{[i+1]}, \ldots, \mathbf{c}^{(j)}_{[i+1]}, \mathbf{c}^{(j+1)}_{[i]}, \ldots, \mathbf{c}^{(N_c)}_{[i]})$.
3. Put $i := i + 1$. If more iterations are required, then go to 2.

Hence the approximations are corrected per species and simultaneously over the grid, such that the diffusion term is treated implicitly. This requires the tridiagonal matrix calculations $2b, 2c$ any time a species is corrected. No Jacobian matrices for the chemistry system are computed and no additional storage is required.

This particular Gauss-Seidel iteration is applied with a fixed number of iterations. Usually only a few iterations already lead to an efficient process. In the tests of this paper, we in fact use only 2 iterations throughout. Although in general we then do not expect to have come very close to the implicit BDF solution, our experience is that using only a few iterations leads to a stable and efficient solution process for tropospheric chemistry models. Comparisons in [VB96, VBH96] with the usual approach based on modified Newton iteration illustrate this. The number of iterations to choose is of course problem dependent. While in [VB96, VBH96] 4 iterations were advocated, a few trial and error runs showed that in the present case 2 iterations will do (see also Section 5.3.6).

5.3.4 The chemistry computation used in type I splitting

Would there be no vertical diffusion ($K = 0$), then the above chemistry-diffusion computation reduces to the chemistry computation used for the box models in [Ver94, VBvLS96]. The same integration formulas are used, the Gauss-Seidel technique is the same and the method is also applied with variable step sizes in the same way as in the chemistry-diffusion computation. We can copy this box model process for the chemistry computation in our type I splitting, since in this splitting we only encounter box models, one at each spatial grid point. The step size for this process will be denoted by Δt_{che} instead of Δt_{cvd}.

5.3.5 The diffusion computation used in type I splitting

On the other hand, without chemistry, the chemistry-diffusion computation reduces to a standard tridiagonal, linearly implicit diffusion integration with the BDF2 formula, one for each vertical column and each species. The Gauss-Seidel process is then no longer operational. In the tests we will take constant step sizes rather than variable ones. Variable step sizes are redundant here. The step size in the diffusion computation will be denoted by Δt_{dif} and taken equal to $\Delta t_{\text{split}}/4$. The extra factor of 0.5 compared to Δt_{adv} is needed because the integration formula is a two-step one.

5.3.6 Mass balance correction

The advection scheme, the implicit BDF integration formula and the spatial discretization of the vertical diffusion all guarantee conservation of mass. However, Gauss-Seidel iteration is not a mass conservative iteration process and since we approximate the implicit BDF solutions with only two iterations, the mass balance can become perturbed. A simple remedy to enforce conservation is to evaluate mass law expressions after each Gauss-Seidel iteration and to correct all species concentrations involved with the ratio (total mass before)/(total mass after). This correction is very cheap with regard to CPU time and can enhance the accuracy notably when the integration is carried out with large step sizes. In the type I splitting procedure the correction is carried out for each of the grid cells and in the type II procedure for each of the vertical columns. The chemistry model introduced in Section 5.4.1 conserves nitrogen. We therefore have applied the above mentioned mass balance correction for all nitrogen compounds.

5.3.7 Reducing the splitting error

At this stage of development we have discussed the full discretization for two operator splitting procedures. Operator splitting is attractive, but it induces a splitting error on top of all the discretization errors. By using Strang splitting, one achieves a reduction of this error, compared to that of the most simple splitting procedure. For pure advection-reaction equations the splitting error can even be eliminated [LY90, HV95]. Spee [Spe95] has implemented this elimination for a 2D advection-reaction problem which is related to the benchmark problem presented in Section 5.4. Thereupon we have used the elimination idea in the tests with the type II splitting procedure reported in the preprint [SdZV+96]. However, in 3D it is theoretically justified only if the wind field is purely horizontal and constant in the vertical direction. Needless to say that this imposes a restriction in practice. We therefore have not implemented the elimination in the tests presented here. A second reason to abandon it is that it would interfere with the comparison between the two splitting procedures, as the elimination is not applicable to the type I splitting procedure.

5.4 A benchmark problem for 3D global transport and chemistry

To test the numerical schemes, to compare the two splitting procedures, and to measure vectorization/parallelization performances, we have developed a benchmark problem for 3D global transport and chemistry. This benchmark problem is similar to the regional test problem from [VBH96] and when we omit the vertical diffusion, it is similar to the 2D problem from [Spe95]. The complete description and a reference solution are available through World Wide Web:
`http://www.cwi.nl/ftp/edwins/Ref_Sol_Benchmark_Global.html`.
To avoid confusion we note that in a number of respects the problem slightly differs from the ones used in the preprints [SdZV+96, SVdZ+97b].

The problem contains horizontal advection, vertical diffusion and chemical reactions. Hence there is no vertical advection and no horizontal diffusion. Orography is not present either, nor do we use real meteorological data so that we are not confronted with massive I/O operations. We realize that this imposes a restriction. However, this restriction enables us to compute a very accurate reference solution which can be used for assessing numerical accuracy and efficiency.

5.4.1 The chemistry model

The chemistry model consists of 45 reactions between 17 species and is used in actual long term global studies where it is referred to as methane chemistry. We obtained it from The [The94]. This model is fully described in the appendix of the preprint to [VBvLS96]. The reaction set gives rise to stiffness. The eigenvalues of the Jacobian matrix lie approximately between -10^9 and 0 [s^{-1}]. There are two extremely large eigenvalues which originate from the free radicals $O(^1D)$ and $O(^3P)$. The chemistry is photochemical. Hence part of the reaction coefficients depend on the solar zenith angle, which depends on the time of the day and the location on earth. Part of the reaction coefficients also depend on the temperature and the pressure. This dependence is chosen in close accordance with the US Standard Atmosphere (1976). We have NO emission at ground level with a constant rate of 10^4 molec/(cm^3 s). The chemistry model is representative for what is currently used in global modeling, although there is a trend towards more complicated chemistry involving more species. All species are advected in the wind field and all species undergo vertical turbulent transport.

5.4.2 Initial condition and time interval

The unit of concentration is number of molec/cm^3 and the unit of time is s. Integration takes place over exactly 14 days, starting at midnight Greenwich Mean Time. With the exception of HNO_3 and NO, for all species the initial concentrations are

at a uniform background level over the globe. Consequently, they are far from the diurnal photochemical equilibrium so that at the start of the integration strong initial transients will be encountered. The ground level values are given in Table 5.1 and initial concentrations in the other vertical levels are such that in each vertical column the mixing ratio (concentration/density) is constant. The initial concentrations of HNO_3 and NO are cylinder shaped, see Figure 5.2. Needless to say that a cylinder shaped initial condition is a challenge for any advection scheme. At the initial time, the center of the cylinder lies at $(\lambda', \phi' = 0°, 0°)$. For HNO_3 and NO, the background concentration is 2.55×10^9 and 100, respectively, and the maximum concentration is 4×10^9 and 10^9. Also for these two species the mixing ratio is taken constant over the vertical layers.

	Name	Species	Concentration
	-	M	2.55×10^{19}
	Water	H_2O	2.55×10^{17}
	Carbon monoxide	CO	2.55×10^{12}
	Molecular oxigen	O_2	5.3295×10^{18}
1	Nitric acid	HNO_3	see main text
2	Peroxynitric acid	HO_2NO_2	1.0×10^2
3	Nitrous acid	HNO_2	1.0×10^2
4	Hydroperoxide	H_2O_2	1.0×10^2
5	Ozone	O_3	7.65×10^{11}
6	Formaldehyde	HCHO	1.0×10^2
7	Methylhydroperoxide	CH_3OOH	1.0×10^2
8	Methylperoxy radical	CH_3O_2	1.0×10^2
9	Methane	CH_4	4.335×10^{13}
10	Nitrogen oxide	NO	see main text
11	Nitrogen dioxide	NO_2	5.1×10^9
12	Nitrate radical	NO_3	1.0×10^2
13	Hydroxyl radical	OH	1.0×10^2
14	Hydroperoxyde radical	HO_2	1.0×10^2
15	Dinitrogen pentoxide	N_2O_5	1.0×10^2
16	Atomic oxigen	$O(^1D)$	0.0
17	Atomic oxigen (g.s.)	$O(^3P)$	0.0

Table 5.1: Initial concentrations at ground level in [molec/cm^3].

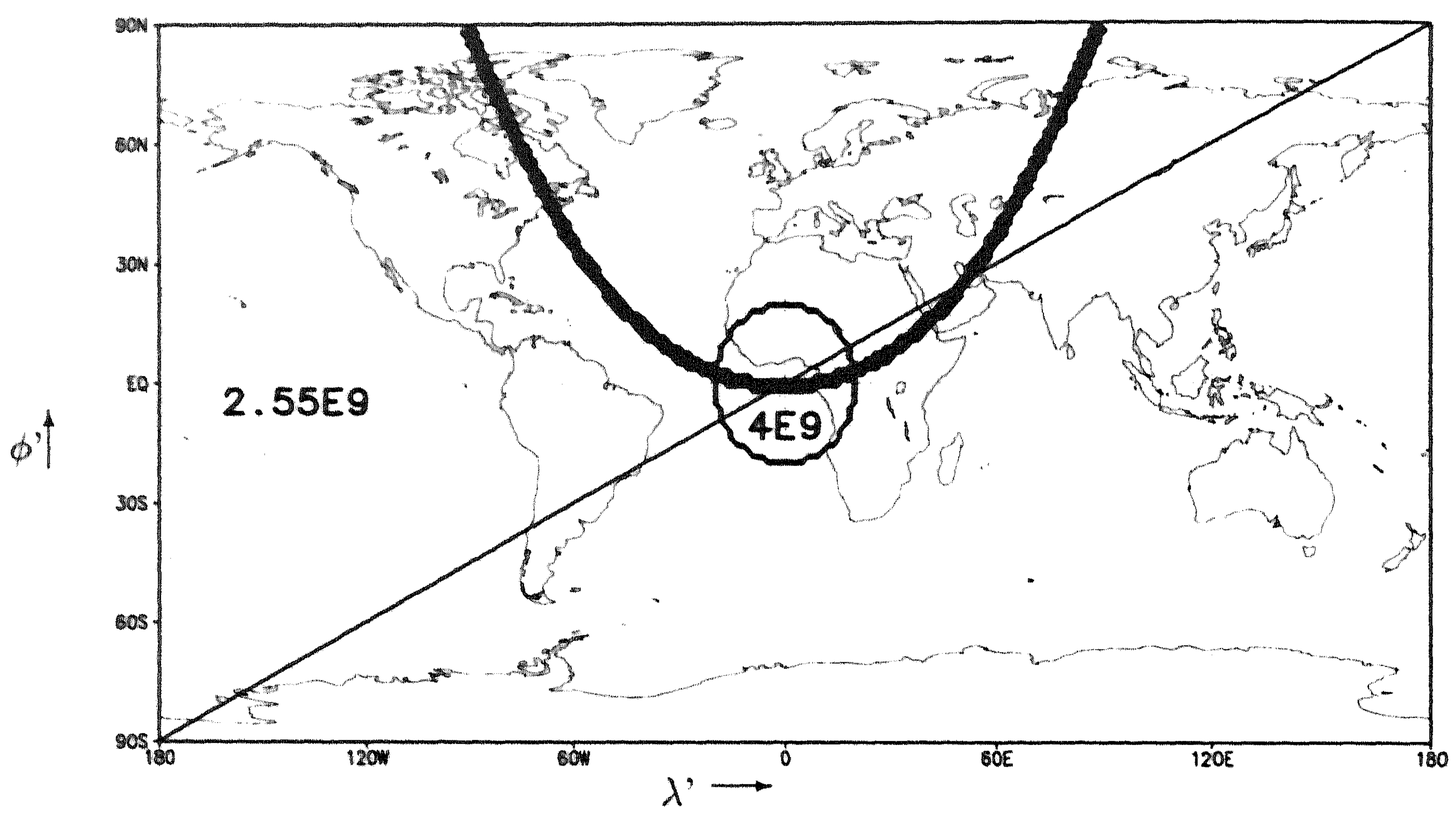

Figure 5.2: The initial concentration vector for HNO_3 and the trajectory of the cylinder. The coordinates are in degrees: $\lambda' = (\lambda - \pi)\frac{180}{\pi}$ and $\phi' = \phi\frac{180}{\pi}$.

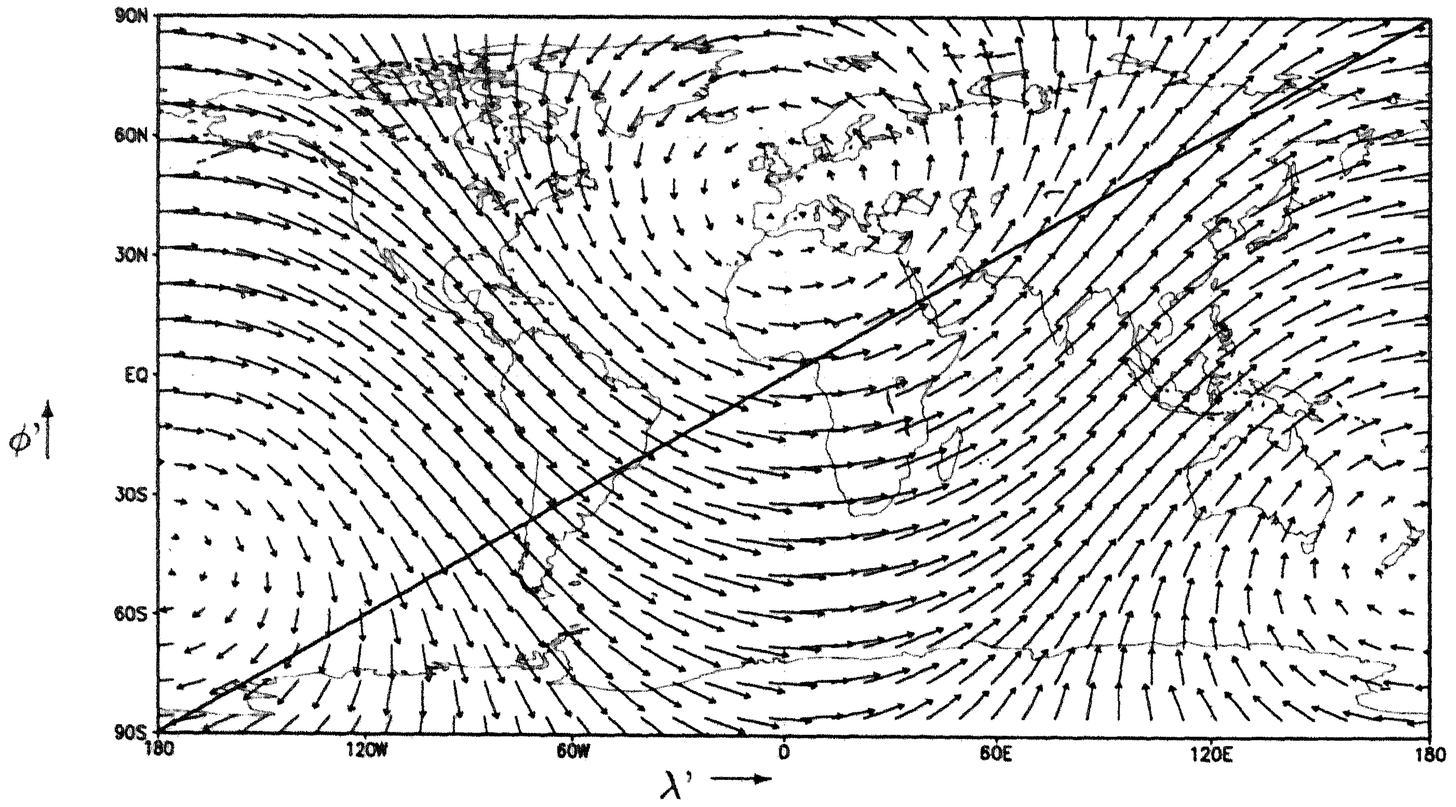

Figure 5.3: The prescribed wind-field ($\lambda' = (\lambda - \pi)\frac{180}{\pi}$; $\phi' = \phi\frac{180}{\pi}$).

5.4.3 Advection

A divergence free, horizontal wind field in analytical form is used to enable the computation of an exact reference solution along characteristics, similar as in [Spe95, VBH96]. Hence vertical advection does not take place. The wind field stems from Smolarkiewicz and Rasch [SR91], and describes a solid body rotation which takes exactly 14 days, see Figures 5.2 and 5.3. The velocities are given by

$$u = 2\pi\,\kappa\,(\cos\beta\cos\phi + \sin\beta\sin\phi\cos\lambda), \quad v = -2\pi\,\kappa\,\sin\beta\sin\lambda, \tag{5.30}$$

where $\beta = 45°$ and $\kappa = a/(14 \times 24 \times 3600)$. Observe that the velocities are constant in the vertical direction and in time. The maximal velocity is approximately 125 km/hour. Near the surface this wind is of course very strong, but higher up in the atmosphere these velocities do occur.

5.4.4 Vertical diffusion

For K have used the K_z depicted in Figure 1-10 in [War88], and approximated as follows (K in m^2/s, z in km):

$$\begin{array}{ll} K = 30 & \text{if } z \le 15 \\ K = 0.2 & \text{if } 15 < z \le 17.5 \\ K = 0.2 + 0.32\ (z - 17.5) & \text{if } 17.5 < z \le 20 \\ K = 10^{\,0.05\ z - 1.0} & \text{if } z > 20 \end{array} \tag{5.31}$$

We have used 15 cell-centered layers to approximate the vertical diffusion term. Simulating a pressure-based vertical grid, the distribution of the cell-centers is a function of the pressure which is taken uniform over the globe. The lowest cell boundary lies at sea level (1000 hPa) and the highest at 38.2 km (0 hPa). The complete distribution of the cell centers, cf. (5.20), reads 0.3, 1.0, 2.2, 4.3, 6.5, 8.4, 10.0, 11.3, 13.0, 15.2, 17.6, 19.8, 22.5, 27.6, 34.7 km.

A characteristic value to assess the numerical stability of a time stepping method for our diffusion problem, is the dimensionless product of the step size Δt and the spectral radius, denoted by σ, of the matrix $\mathbf{A}$ introduced in Equation (5.24). An upper bound for σ is found by applying Gerschgorin's theorem, which in the present situation yields

$$\Delta t\sigma \approx \Delta t \max \frac{4K}{(\Delta r)^2}, \tag{5.32}$$

where the coefficient K and the grid distance Δr vary over the grid. If $\Delta t\sigma \approx 1$, then explicit time stepping will do for numerical stability. If this product is $\gg 1$, however, an implicit approach is necessary. Our current choice for the grid and the coefficient K yields

$$\Delta t\sigma \approx \Delta t\,\frac{4 \times 30}{700^2} = \frac{\Delta t}{4100}, \tag{5.33}$$

which is fairly small in the sense that even a step size of approximately 1 hour still allows an explicit treatment. Obviously, if in the boundary layer a much finer grid would be used, like in regional and urban models, then the stability restriction can be much more severe and an implicit treatment is necessary.

In both splitting procedures we therefore treat the diffusion term implicitly as outlined in Section 5.3.3. Since in our approach this gives rise to the solution of tridiagonal linear systems only, the additional costs compared to an explicit treatment are not high. For other methods solving chemistry and vertical diffusion in a coupled manner, the above considerations can have a greater impact for the efficiency of treating vertical diffusion explicitly or implicitly. Needless to say that the range of values K takes on should also be taken into consideration. Our maximum value of 30 m^2s^{-1} seems rather high, compared to for example the maximum value of 10 used in [GM91]. In other words, it may well be that in actual global models, where the vertical grid sizes in the planetary boundary layer will not be smaller than 100 m, say, and the maximum for K is not greater than 10 m^2s^{-1}, an explicit treatment can be used without a severe time step restriction.

5.4.5 The reference solution

In case of pure advection, the cylindrically shaped initial profiles for HNO_3 and NO would be transported over the northern hemisphere as depicted in Figure 5.2 and return at their initial position at the diagonal line $\lambda' = \phi'/2$ in exactly 14 days. We therefore have computed an accurate reference solution at all vertical layers lying above this whole line. Obviously, vertical diffusion and chemistry change the solution profile. But still the test is quite meaningful for advection schemes.

Without vertical diffusion, computing an accurate reference solution can be done by backward solving the chemistry along characteristics. To circumvent the difficulty of vertical diffusion, which prevents this, we have followed this procedure for the semi-discrete system which results from discretizing the diffusion operator, similar as in [VBH96]. Hence our reference solution can be considered exact except for the error due to the spatial discretization of the vertical diffusion term. This means that assessment of accuracy is possible only for the horizontal advective transport, the solution of the chemistry and the time integration of the vertical diffusive transport.

Figure 5.4 shows the reference solution for O_3, NO_X, HNO_3 and HO_2NO_2 along the diagonal line at ground level. One can see that a rather fine grid is needed to resolve all details and that as a result of the chemical reactions a cylindrical profile is also present for species other than HNO_3 and NO.

5.4.6 Problem size

To illustrate the scale of the numerical problem, we included Table 5.2 which for three practical choices of the spatial grid gives the total number of concentration values that must be computed by integration in time. In this table we have taken into account the

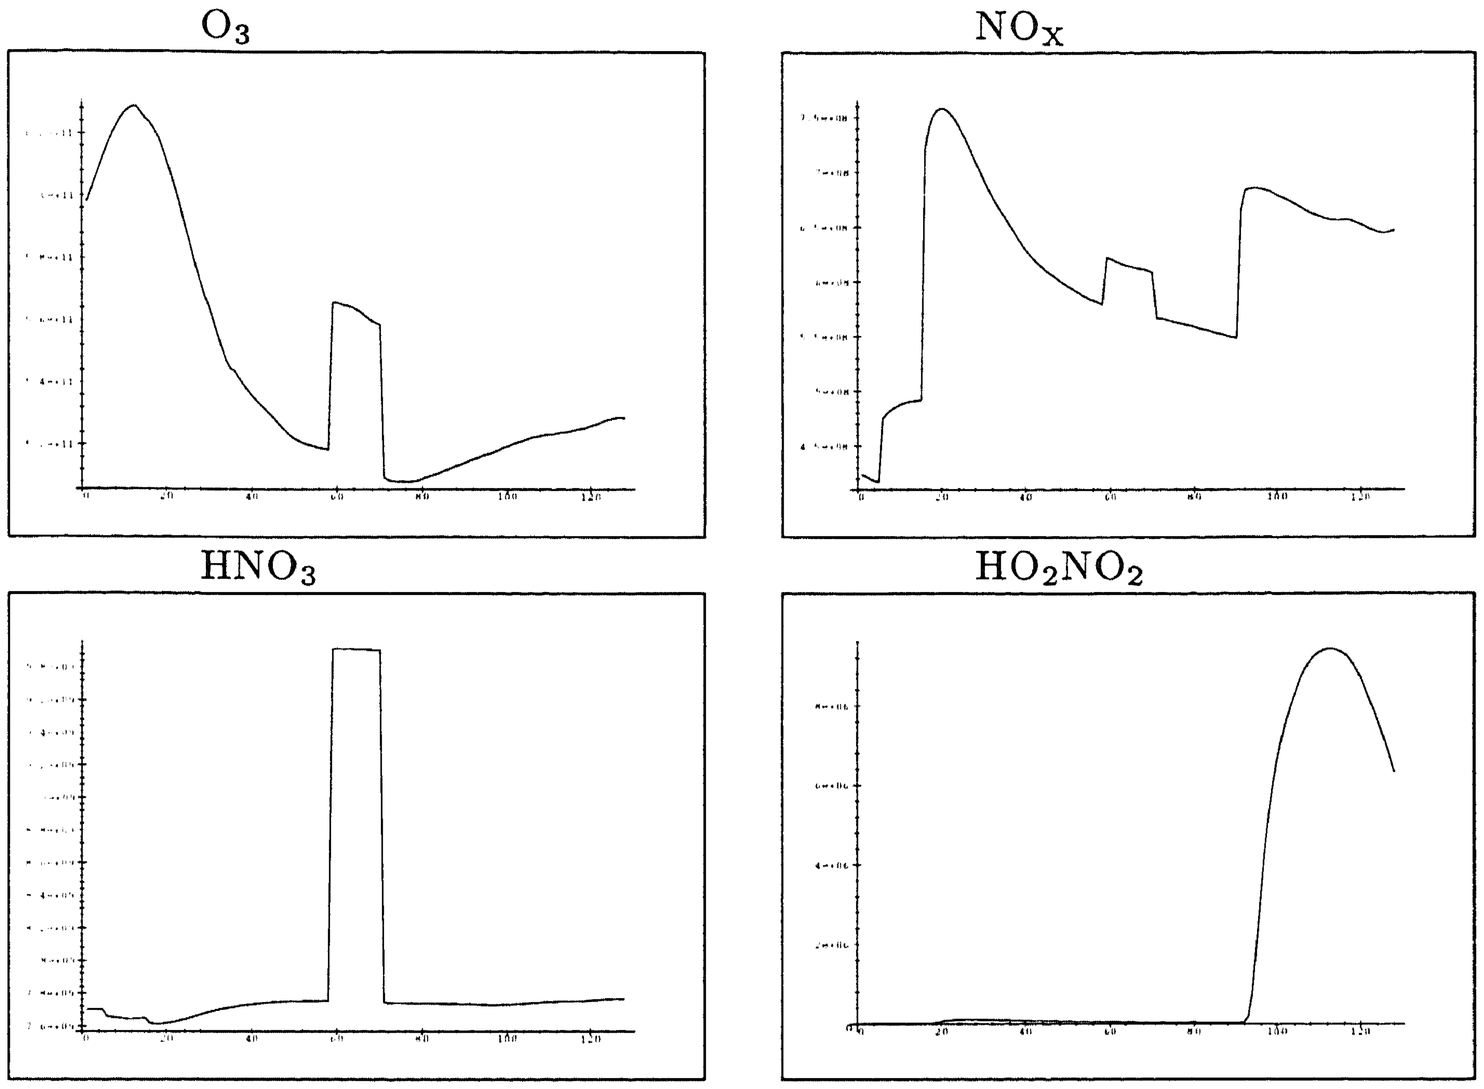

Figure 5.4: Reference solutions at ground level of O_3, NO_X, HNO_3 and HO_2NO_2 in unit molec/cm^3 (vertical axis) plotted along the diagonal $\lambda' = \phi'/2$ (horizontal axis).

savings of the grid reduction near the poles. The total number of unknowns clearly illustrates that global air quality modeling leads to a huge computational task. In this connection it is important to recall that our time span of 14 days is short and that our chemical scheme composed of 17 species is of moderate size, see [PBC+95]. The same can be said about the space grids. The finest longitude/latitude grid here is approximately $1.4° \times 1.4°$. In [PBC+95] it is pointed out that still finer grids are needed in future generation models.

Table 5.2: Problem size.

Grid (λ, ϕ) (r=15)	uniform grid cells	unknowns on uniform grid	grid reductions	reduced grid cells	unknowns on reduced grid
64×32	30 720	522 240	4	24 840	422 280
128×64	122 880	2 088 960	5	93 960	1 597 320
256×128	491 520	8 355 840	6	391 560	6 656 520

5.5 Vectorization, performance, accuracy and efficiency

Our tests have been carried out on a Cray C90. It thus is of interest to optimize the code with respect to vectorization. In this section we outline the vectorization, we present results of performance tests on the benchmark problem and comment on the achieved numerical accuracy and efficiency.

5.5.1 Vectorization

On a uniform grid, vectorization of the explicit advection computation is straightforward and can easily be optimized. In [BHV94] it is shown that the non-uniformity near the poles only leads to a minor reduction of the vector speed, mainly because the larger part of the grid is still uniform. However, even on the non-uniform part of the grid the vector speed is still quite acceptable. For the technicalities and modifications required near the poles we refer to [BHV94].

The vectorization of the coupled chemistry-diffusion computation in the type II splitting procedure is similar to that in [VBH96] and is carried out along the horizontal grid dimensions. In a different setting, this idea of vectorization was first proposed in [GvL90]. As far as we know, in air quality modeling it was first reported in [JT94]. Here it works because the chemistry-diffusion computation results in a coupling only in the vertical direction, which means that all operations involved can be executed simultaneously over the horizontal grid points (λ_i, ϕ_j).

However, as pointed out in [JT94], there is a snag in the choice of the step size Δt_{cvd}. For grid-vectorization we would like to use one and the same integration step size over the horizontal grid, because then the vector length would always be sufficiently large for reaching the optimal vector speed. Unfortunately, this would cause a reduction of efficiency, because at parts of the globe we have to obey step size restrictions more severe than elsewhere, viz. in the (constantly changing) areas on the globe where day turns into night and vice versa. In these areas the species undergo intense chemical reactions since the chemistry is photochemical. This gives rise to rapid changes in concentrations. On the other hand, during nightly periods and also during day time, changes are much slower, so that adjusting the step size to the sunset/sunrise situation everywhere, is less efficient. As a compromise, we therefore group points (λ_i, ϕ_j) in clusters of length 128, the vector length of the C90. This grouping of points into clusters is in the ϕ direction, in order to stay in the same time zone as much as possible. The chemistry-diffusion computation is then performed and vectorized for a single cluster using the same step size Δt_{cvd} for this cluster, while step sizes may differ per cluster. This compromise works out satisfactorily, although we will still encounter step sizes that are too small for part of the grid points. We emphasize that our clustering technique differs from the one used in [JT94]. In [VBH96] clustering has not been considered.

The vectorization of the diffusion computation in the type I splitting procedure goes entirely similar, i.e., also along the horizontal grid dimension. But the clustering issue is of no relevance here, as constant step sizes Δt_{dif} are used. The chemistry computation does use clustering though, since the step size Δt_{che} is variable. Because only box models appear, the clustering can now be carried out also in the vertical direction so that the clustering region in the horizontal plane can be kept much smaller.

It is also possible to do part of the computational work simultaneously over the vertical direction as well, e.g. the computation of the production terms, which offers a way of increasing the vector length. In order to realize this profit we implemented separate subroutines for performing essentially the same computations for the uniform and the non-uniform part of the grid. So-called (Cray) CDIR-directives were not needed because we could benefit from the loop collapsing (i.e. merging nested loops into one single loop) done by the (Cray) `FORTRAN` preprocessor (FPP).

5.5.2 Performance results

We ran our code, written in `FORTRAN 77`, on the Cray C90 at SARA, Amsterdam. Timing results presented in this section were done on one processor with a clock cycle time of 4.2 nanosecond and a double vector pipe. This gives a theoretical peak performance on one processor of 476 Mflop per second and 952 when chaining an add and a multiply. Chaining floating point operations, however, can only be achieved for very specific loops so that in practice a performance of 50% of the peak performance (500 Mflops) can already be considered as a very good result. To measure the Mflop

rate and the CPU time of a routine, we used the Cray utility `perftrace` that gives the hardware performance by program unit.

The discretization parameters were chosen as follows. Three different reduced longitude-latitude grids were used with the same vertical grid, see Table 5.2. These grids determine the critical advection step size Δt_{adv} for stability and positivity through the CFL condition (5.17). The critical values are, approximately, 40, 20 and 10 min. In our tests these values determine $\Delta t_{split} = 2\ \Delta t_{adv}$ and $\Delta t_{dif} = \Delta t_{adv}/2$. The step sizes Δt_{che} in splitting I and Δt_{cvd} in splitting II vary in time and are governed by the step size control mechanism of the integration method. Both are constrained by a minimum of 5 min throughout. This minimum step size was selected after a few trial and error runs. This minimum value is important as it determines to a great extent the CPU time and the accuracy of the chemistry integration.

Table 5.3 shows performance results for the two splitting procedures on the three different grids. The values for Δt_{cvd} and Δt_{che} given in this table are average values. For the two procedures as a whole we do not observe large performance differences. Therefore we will only comment on the type II procedure.

The overall Mflop rate is close to 500 which means close to 50% of the peak formance. On the coarsest grid the Mflop rate of 371 in the advection computation is low due to the grid reduction. The finer the grid, the smaller the influence of grid reduction will be so that larger flop rates are obtained. When comparing CPU times needed for the advection and the chemistry-diffusion computation, we see that on the coarsest grid the latter is three times more expensive, while on the finest grid the advection takes 46% of the total CPU time versus 54% for the chemistry-diffusion computation. The advection becomes more costly due to the diminishing step size Δt_{adv}, whereas the step size Δt_{cvd} of the chemistry-diffusion computation hardly changes and stays close to the imposed minimum of 300 s. Three reasons for this behavior can be mentioned. First, we always start with the minimum of 300 s which determines a certain maximum value. For example, on the finest grid, where the split interval is only 1200 s long, the smallest possible number of steps with our two-step method equals three so that the largest possible average for Δt_{cvd} is only 400 s (= (300 + 300 + 600)/3 s). Of course, on the coarsest grid there is more time for Δt_{cvd} to increase. Second, the initial transients introduced at the beginning of each split interval do withhold the step size from increasing very rapidly directly from start on. As the third reason we recall the global nature of the problem which means that we always have to integrate through sunsets and sunrises during which rapid temporal changes in concentration values exist.

The imposed tolerance for the step size control also plays a role in the above. We have not examined this issue further, since a step size of $\approx$ 5 min for a low cost chemistry-diffusion computation is already quite efficient, in our opinion. It is obvious, though, that we might as well have used a constant step size of 5 min throughout which makes the cluster approach redundant and also avoids the costs of the step size control.

Table 5.3: Performance on one processor of the C90.

(a) Results for the type I splitting procedure.

Grid (λ, ϕ, r)	Advection				Diffusion				Chemistry			
	Δt_{adv} (s)	CPU %	rate (Mflops)		Δt_{dif} (s)	CPU %	rate (Mflops)		Δt_{che} (s)	CPU %	rate (Mflops)	
64×32×15	2400	30	366		1200	5	331		313	62	416	
128×64×15	1200	35	482		600	12	225		312	53	410	
256×128×15	600	42	525		300	16	224		300	43	391	

Grid (λ, ϕ, r)	Total	
	rate (Mflops)	CPU (s)
64×32×15	384	485
128×64×15	412	2330
256×128×15	420	14700

(b) Results for the type II splitting procedure.

Grid (λ, ϕ, r)	Advection			Chemistry-diffusion		
	Δt_{adv} (s)	CPU %	rate (Mflops)	Δt_{cvd} (s)	CPU %	rate (Mflops)
64×32×15	2400	28	371	301	71	524
128×64×15	1200	37	487	322	63	498
256×128×15	600	46	529	311	54	451

Grid (λ, ϕ, r)	Total	
	rate (Mflops)	CPU (s)
64×32×15	477	496
128×64×15	493	2190
256×128×15	486	13300

5.5.3 Accuracy

For the six tests tabulated in Table 5.3, Figures 5.5 - 5.8 show plots of the computed and reference solution concentrations of O_3, NO_X, HNO_3 and HO_2NO_2 in the same way the reference solutions were depicted in Figure 5.4.

For the chosen discretization parameters, both splitting procedures can be seen to work very well. Even the coarsest grid results are already reasonably accurate for the actual practice, while going to a finer grid clearly resolves all details in the solutions. Noteworthy is the excellent performance of the advection scheme in resolving the cylindrical profiles. In particular the accuracy for O_3 is excellent. The accuracy for HNO_3 lags behind. This is due to the fact that we use only two Gauss-Seidel iterations in the chemistry integration. Resolving the chemistry more accurately will reduce the errors for HNO_3.

Of interest is that decoupling stiff chemistry and vertical diffusion does not lead to notably larger errors. The coupled procedure (type II splitting) is more accurate, but the differences are not really significant. In this respect our findings seem to contradict those from [GM91] and [SVdZ+97b], where large splitting errors resulting from decoupling stiff chemistry and vertical diffusion are reported. We owe this to the difference in K-values. In [SVdZ+97b] a K was used 10^6 times larger than in the current benchmark. In [GM91] the K-values are of the same magnitude as here, but the height extends to 1 km only versus 38.2 km in our benchmark.

5.5.4 Efficiency

Efficiency and the required CPU times are more difficult to assess since we cannot compare our results with those obtained by other solvers. However, the following observation indicates that our splitting method II and its vectorized implementation are quite efficient. Table 5.3 shows that the CPU time for the advection computation on the finest grid is only 12% less than for chemistry and vertical diffusion together. This means that chemistry and diffusion are dealt with efficiently when taking the advection computation as a reference point. Of course, the accuracy should be sufficiently high to render this conclusion of practical value. We have shown that this is indeed true for the current benchmark problem.

It is also of some interest to compare simulation time/CPU time ratios for our benchmark tests with the predictions given in Table 3 of the review paper [PBC+95]. For a hypothetical model, representative for the state of the art, that table predicts required flop rates for three different grids and three different ratio's. For our $1.4° \times 1.4° \times 15$ grid computation, lower table (b), we have a ratio of 85:1 against ≈ 0.5 Gflops. Comparing this ratio with the ones given in [PBC+95], reveals a wide gap, clearly to the advantage of our solver. For example, to achieve a 100:1 ratio for a grid resolution of $2° \times 2° \times 10$, the flop rate prediction in [PBC+95] is 20 to 30 Gflops, whereas we are close to this ratio with ≈ 0.5 Gflops on a finer grid. This again indicates that our solver is efficient. However, care is needed here as there are many uncertainties in this

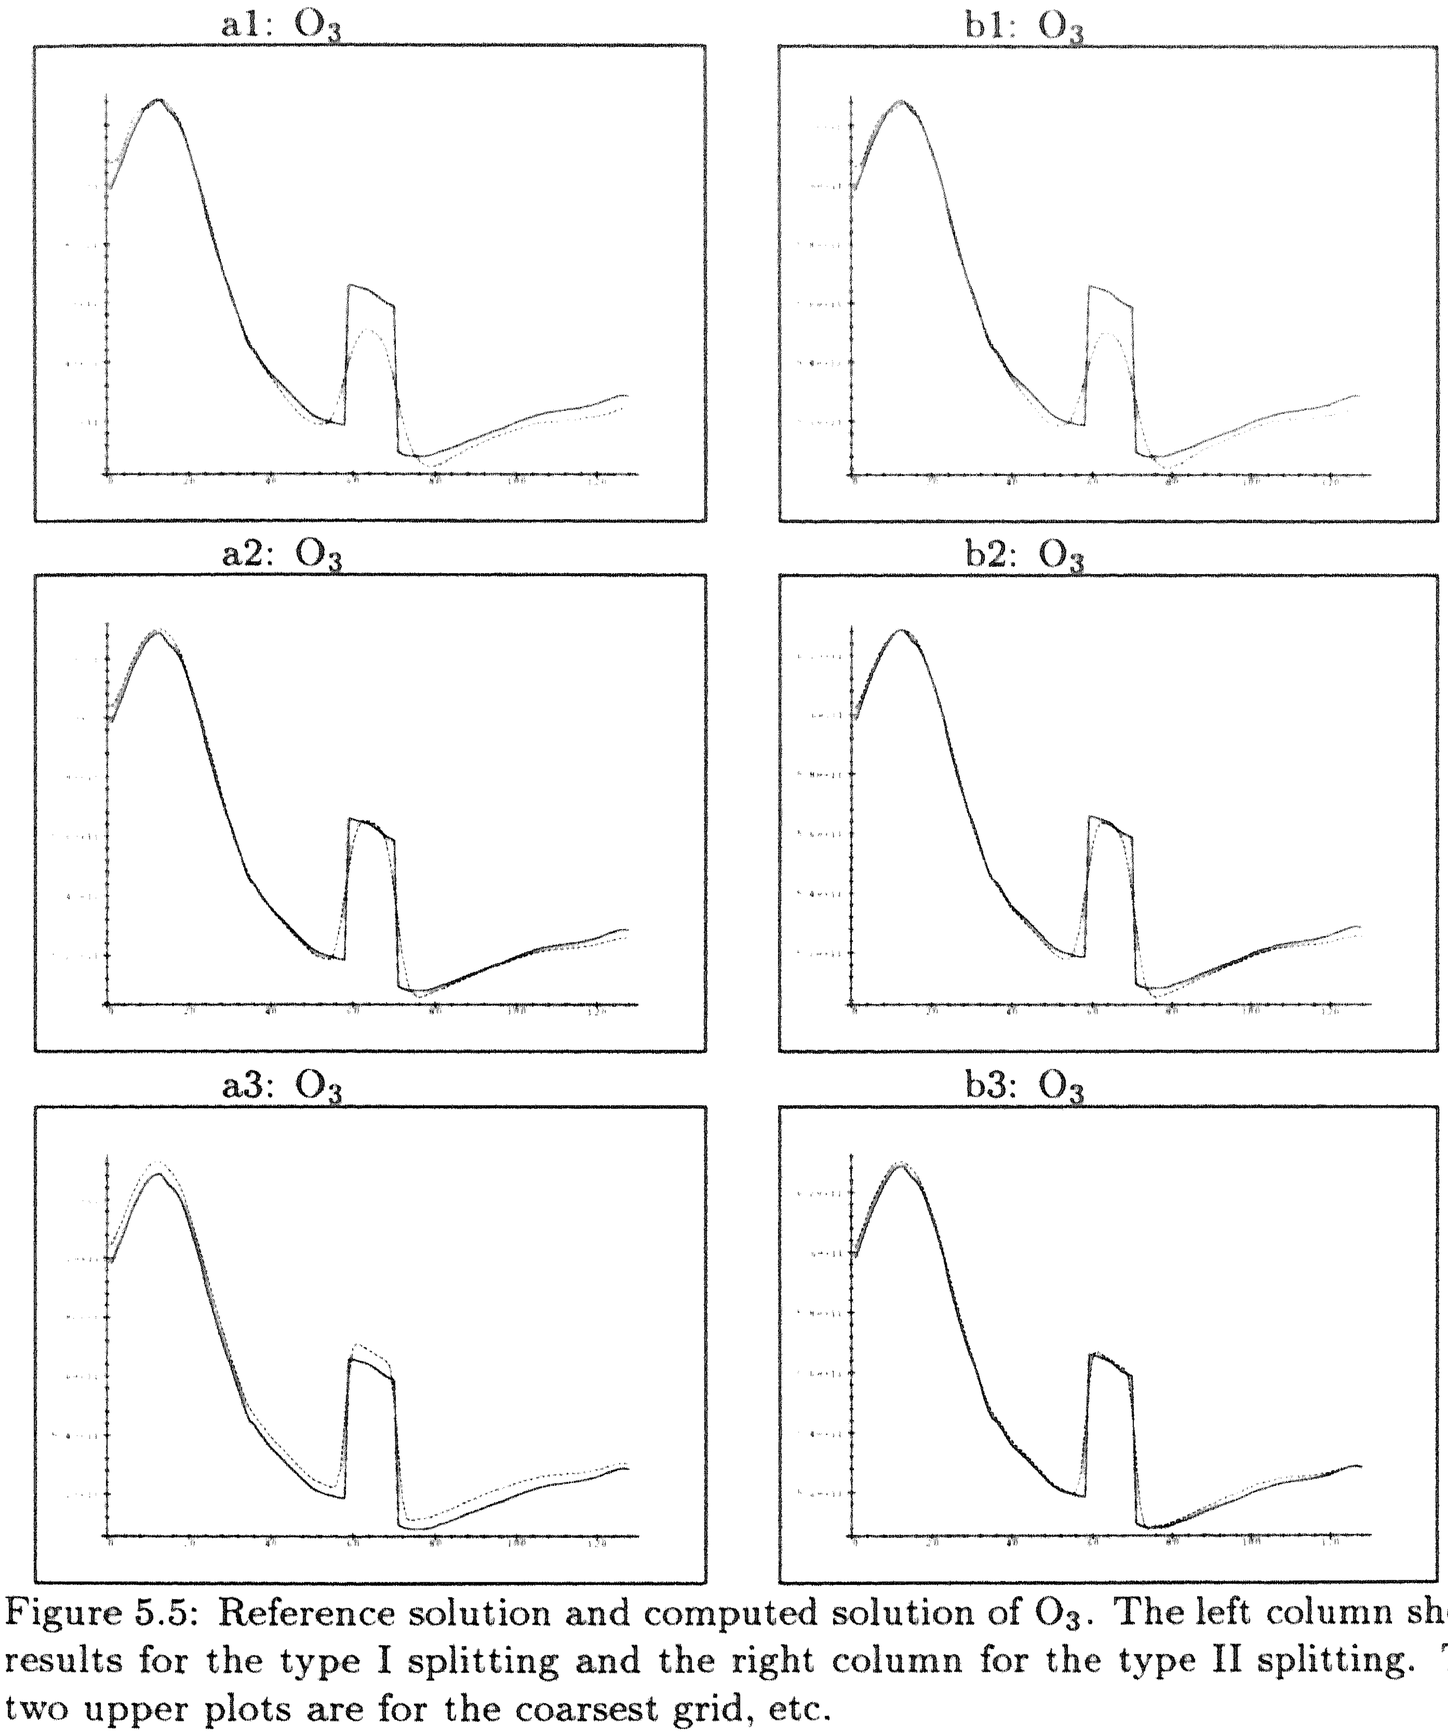

Figure 5.5: Reference solution and computed solution of O_3. The left column shows results for the type I splitting and the right column for the type II splitting. The two upper plots are for the coarsest grid, etc.

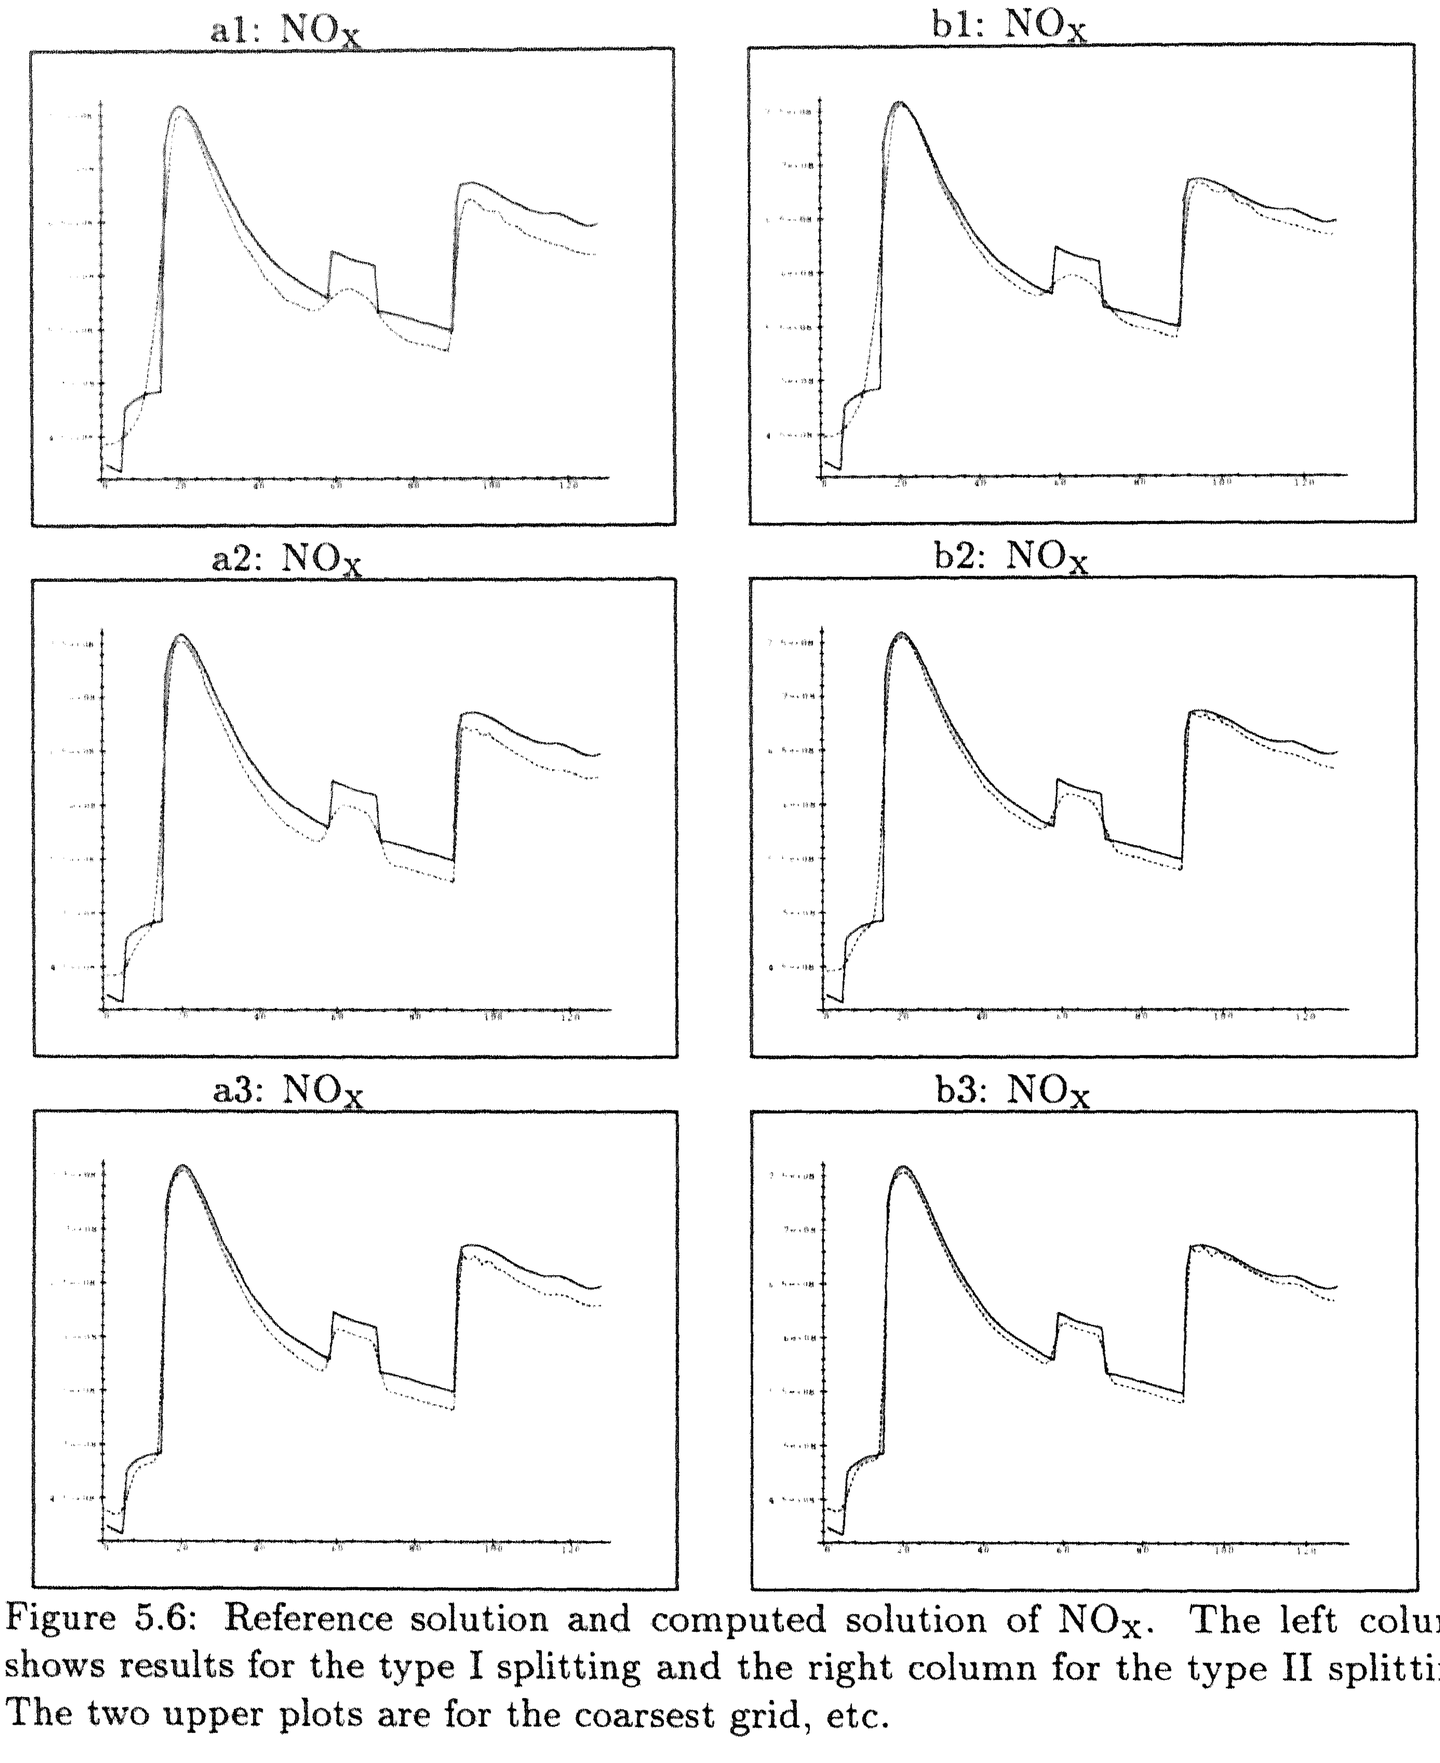

Figure 5.6: Reference solution and computed solution of NO_X. The left column shows results for the type I splitting and the right column for the type II splitting. The two upper plots are for the coarsest grid, etc.

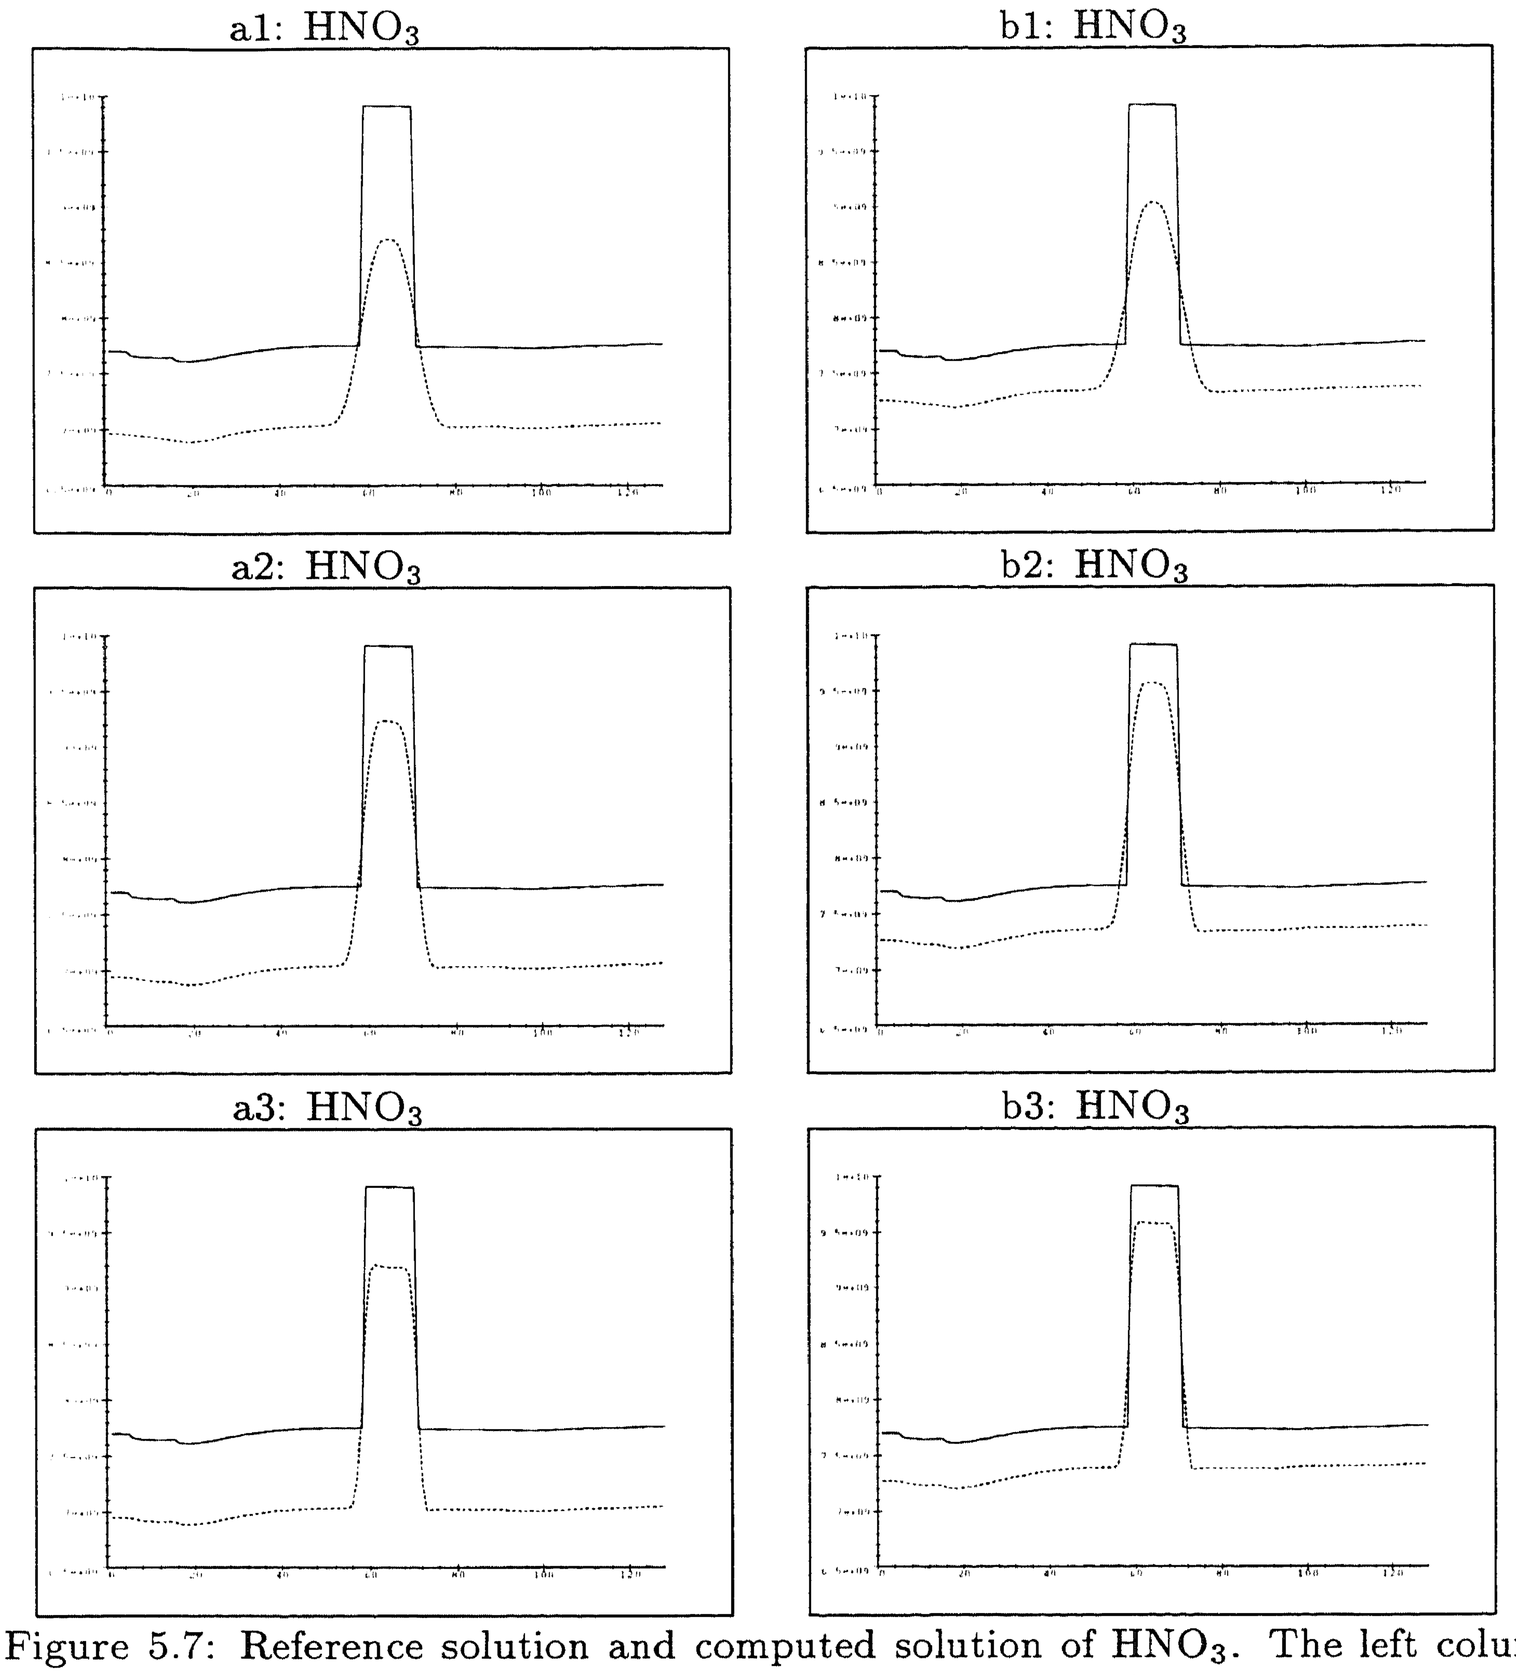

Figure 5.7: Reference solution and computed solution of HNO_3. The left column shows results for the type I splitting and the right column for the type II splitting. The two upper plots are for the coarsest grid, etc.

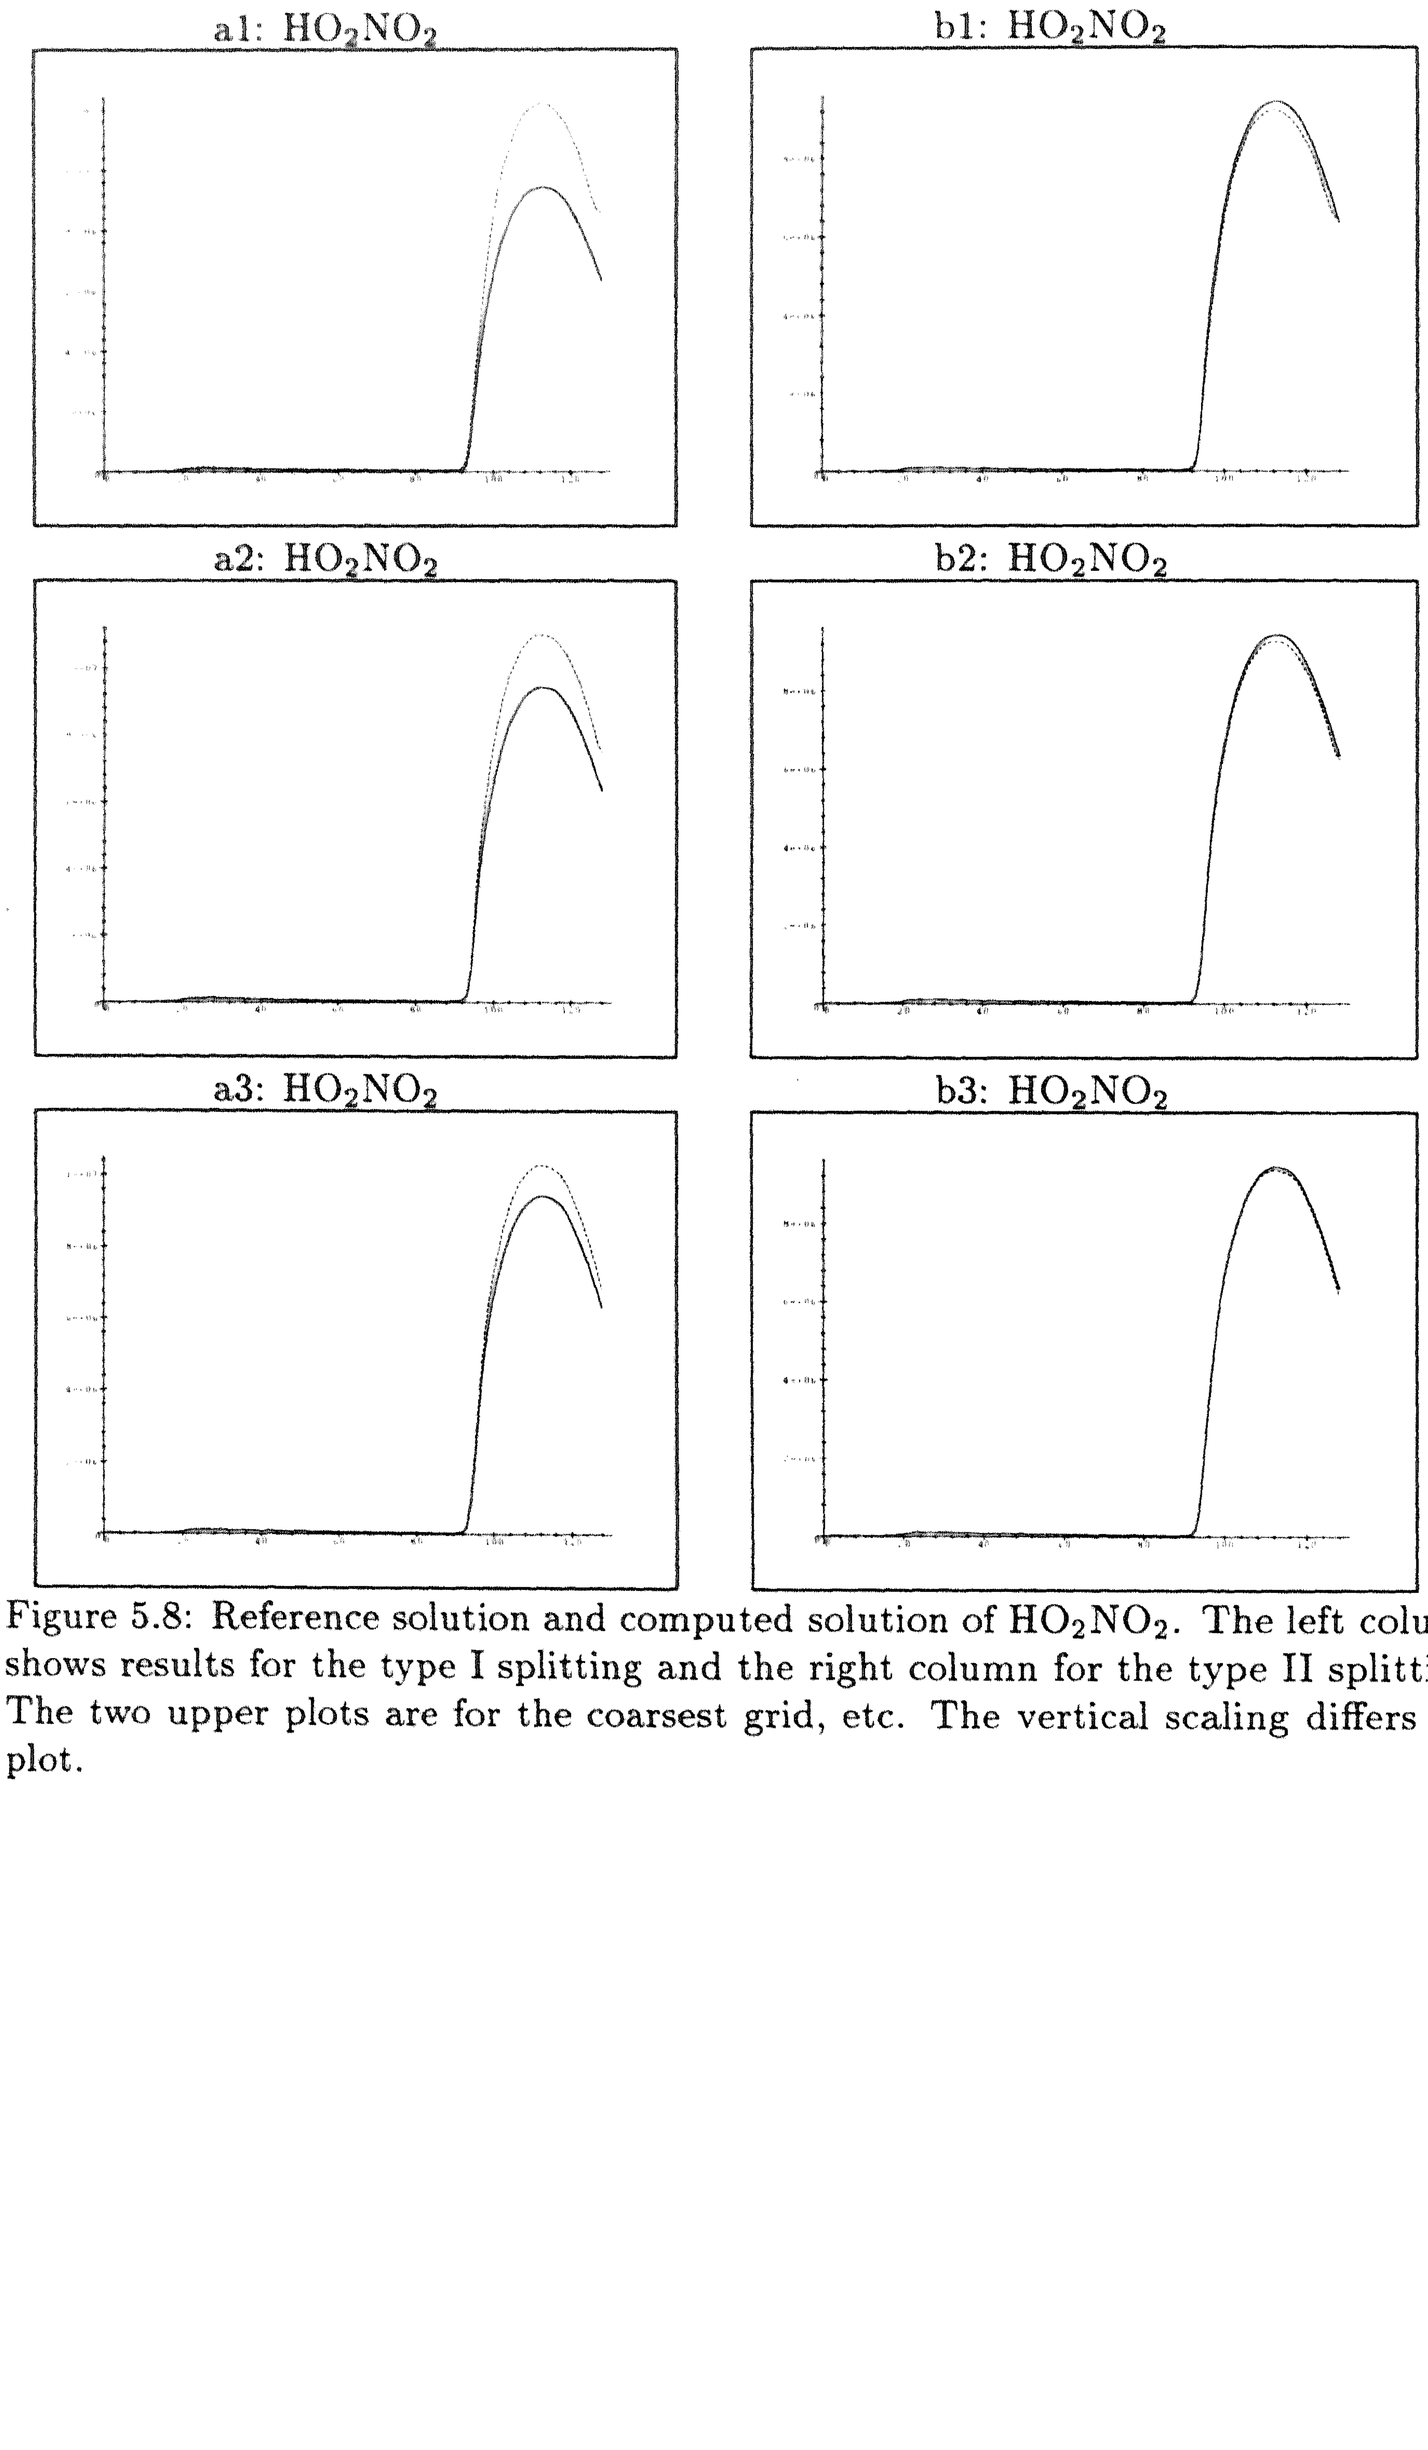

Figure 5.8: Reference solution and computed solution of HO_2NO_2. The left column shows results for the type I splitting and the right column for the type II splitting. The two upper plots are for the coarsest grid, etc. The vertical scaling differs per plot.

comparison, an important one being the computational complexity of our benchmark compared to that of the hypothetical problem of [PBC+95].

5.6 Parallelization on the C90

While vectorization effectively reduces CPU time and takes place on a single processor, parallelization merely reduces the wall-clock time by distributing the work over multiple processors. In this section we discuss the parallelization of the type II splitting procedure, for which two different approaches have been considered, viz. autotasking and a parallelization over the geometry.

Autotasking can be described as the automatic distribution of loop iterations to multiple processors [Cra91]. This type of parallelization takes place on the level of elementary algebraic operations and dependency analysis across procedure boundaries does not take place. Autotasking takes a `FORTRAN` program as input, then modifies it and adds compiler directives, so that it can run concurrently on multiple processors. This works best on programs in which most of the work concerns independent operations in nested do-loops. If possible, the innermost loop of a nest of do-loops is vectorized and autotasking runs the outermost loop on multiple processors. In some cases, autotasking will process a single vectorizable loop into chunks.

Because the advection computation is explicit, for this part we can rely on autotasking (the explicit Runge-Kutta method invokes merely elementary algebraic operations on long vectors). Autotasking of the chemistry-diffusion problem is less efficient, however. For this part we have implemented an alternative which, conceptually, is a simple form of domain decomposition. Recall that for the grid-vectorization we group points (λ_i, ϕ_j) on the globe in clusters of length 128. But now we have p processors available instead of only one. Hence we can distribute clusters over the processors and use, on each processor, the same vectorized chemistry-diffusion computation. This way of parallelization is much more effective than autotasking. So-called (Cray) CMIC-directives were used to assign each cluster to a different processor. By such directives we could parallelize loops containing calls to subroutines which are on a high algorithmic level. One needs to be careful though and indicate clearly which variables are global data and which are local.

Load balancing refers to the even distribution of work over all processors. Here, we need to reckon with a step size Δt_{cvd} that varies from one cluster to another. Therefore, one cluster may take more time steps within a splitting step than another one. This may affect the load balancing. Nevertheless the average effect appears to be moderate, as can be expected since Δt_{cvd} is in practice approximately constant.

Regarding our expectations on parallel speed up, we have to bear in mind the restriction put by Amdahl's law. Suppose we have a number of p processors available for parallel execution of a code. Let $S = T(1)/T(p)$ denote the speed up with $T(p)$ the wall-clock time required to execute the code on p processors. Amdahl's law then

reads

$$S = \frac{1}{(1-f) + f/p}, \tag{5.34}$$

where f denotes the fraction of the work that can be executed in parallel. The effects of parallelization can be measured by means of the Cray tool `atexpert`. This tool predicts speedups on an almost dedicated machine from data collected from a run on a non-dedicated machine [Cra91]. Table 5.4 contains the predicted values for the three test cases (b1) - (b3) of Table 5.3, assuming 2, 4 and 8 processors. The bracketed numbers correspond to the ideal speedup (or parallel fraction f) according to Amdahl's law (5.34). The results are self-evident. For example, on the coarsest grid the actual speedup for 8 processors is predicted to be 4 and on the finest grid 6.6. These expected speedups are very satisfactory when taking into account that the grid dimensions have not been adjusted to the number of processors. The corresponding parallel fraction for these two cases amounts to 93% and 99%, respectively. Finally, the predictions are in accordance with those from [VBH96] (Table 3).

Table 5.4: Speedup factors for type II splitting.

Grid (λ, ϕ, r)	p = 2	p = 4	p = 8
64 × 32 × 15	1.8 (1.9)	3.0 (3.3)	4.0 (5.3)
128 × 64 × 15	1.8 (2.0)	3.2 (3.8)	5.6 (7.0)
256 × 128 × 15	1.9 (2.0)	3.6 (3.9)	6.6 (7.6)

5.7 Conclusions, final remarks and future work

Air quality modeling on a global scale is numerically extremely expensive and requires fast algorithms and sophisticated numerical software on high-speed computers. Using a powerful, shared memory, vector/parallel computer, a Cray C90, we have compared and tested two operator splitting procedures for a problem that we would like to propose as a benchmark for 3D global transport-chemistry solvers. There is a great need for model problems in this field. Without representative large-scale model problems, comprehensive testing is hardly possible, let alone comparison and validation of solvers. The interested reader is invited to join us in this benchmark activity and to apply his own solver to our test problem, see:
`http://www.cwi.nl/ftp/edwins/Ref_Sol_Benchmark_Global.html`.
The present test results justify the following conclusions:

- We advocate to keep the stiff chemistry and vertical turbulent diffusion coupled, as this alleviates the onset of transients and reduces the splitting error for very fast reacting species. For other species, O_3 e.g., decoupling has no adverse effect on the accuracy in the current test model. Yet, in spite of the 1D nature,

a coupled solution can be rather expensive when standard modified Newton iteration is used within an implicit integration approach. Our type II splitting procedure is special in that it solves the stiff chemistry coupled with the vertical diffusion using a tridiagonal, Gauss-Seidel iteration. This iteration technique reduces the costs of the coupled approach significantly and results, together with the chosen advection scheme, in an efficient procedure for tropospheric gas-phase transport-chemistry models.

- Significant effort has been put in vectorizing the algorithms on the C90. We have shown that vectorization along the horizontal grid dimension is close to optimal. On fine grids the overall Mflop rate is about 500. This means that also the explicit horizontal advection scheme, which makes use of a reduced grid near the poles to alleviate the stability restriction, vectorizes very well.

- These high flop rates were measured on one processor. We also examined parallelization of the type II splitting procedure and have shown that by a rather straightforward approach, a very satisfactory speedup is realized. On fine grids, as considered for future practical applications, 75% of the optimal speedup seems within reach on the shared memory, vector/parallel C90 machine.

In the near future our research will be continued into two directions. First, a second benchmark problem will be considered involving real meteorological data and orography. The main purpose hereby is to further test and compare our reduced-grid advection scheme under more realistic conditions. A report on this investigation has already been finished [SPvDH97]. Second, alternative splittings will be examined which reduce the initial transient complications in the chemistry integration. One of our aims hereby is to further benchmark the recently proposed sparse Rosenbrock solvers [SVB+97] for application to transport-chemistry problems.

Chapter 6

Sensitivity of atmospheric transport model performance to numerical advection schemes and resolution

Abstract

The need to use high horizontal resolutions in global transport models in order to adequately model anthropogenic perturbations to global atmospheric chemistry, has led to a demand for cheap advection schemes with high accuracy. We have tested four newly developed 3D advection schemes named Mol-rg, Split-u, Split-us and Split-rg. We compared these schemes with Slopes and Second Moment. Mol-rg and Split-rg make use of a reduced grid, a grid with less cells near the poles, to overcome the well known pole-singularity.

Two tests were performed with all schemes: the solid-body rotation test and a radon test. The radon test uses 3D meteorological input for the month January 1992 from a numerical weather prediction model, together with parametrizations for sub-grid scale vertical transport. We compared model results with measurements on two islands in the Indian Ocean.

The solid-body rotation test showed that none of the new schemes generates under-shoot and overshoot and that all of them are mass conservative. Split-us and Split-rg are very cheap in terms of cpu time and memory requirements and give accurate results.

The model results in the radon test were found to be almost independent of the numerical scheme used. The results depend more on the quality of the analyzed wind fields and sub-grid scale vertical transport that are used as input to the model.

Therefore it is advised to use cheap advection schemes such as Split-rg to be able to perform model calculations on high resolutions.

The new advection schemes are available through Internet.

Note: This chapter is a revision of report MAS-R9710 and has been submitted for publication in *Journal of Geophysical Research*, and is currently under revision.

6.1 Introduction

Due to non-linearity of chemistry and spatial heterogeneity of anthropogenic trace gas sources, a trend towards the use of higher resolutions in global tropospheric chemistry modeling can be recognized. For models designed to run for several years with a large number of chemical species and using monthly averaged climatological meteorology, the horizontal resolution has recently increased from (latitude $\times$ longitude) $10° \times 10°$ in MOGUNTIA [CZ91] to $5° \times 5°$ in IMAGES [MB95]. At this moment both on-line models, which calculate the meteorology at every time step, and off-line models, which use climate or weather forecast model output at 4 to 12-hourly time resolution, typically use a resolution between $3.8° \times 3.8°$ and $5.6° \times 5.6°$ to simulate tropospheric ozone and sulfur species [RLSK97, FKR+96, CJG+96]. According to Peters et al. [PBC+95] grid resolutions of $0.5° \times 0.5°$ or better in the horizontal are needed for an adequate modeling of anthropogenic perturbations to global atmospheric chemistry.

This paper addresses the computational aspects involved in increasing the resolution of present-day on-line and off-line atmospheric chemistry models. We focus on the cost of numerical solution of the advection equation which for most schemes increases more rapidly for higher resolutions than the numerical solution of the chemistry equations. This is due to the CFL constraint on the advection time step. We investigate several aspects of the performance of two widely used advection schemes (Slopes [RL81] and Second Moment [Pra86]) and four new schemes (named Split-rg, Split-u, Split-us and Mol-rg) at horizontal resolutions varying from $5.6° \times 5.6°$ to $1.4° \times 1.4°$. The sensitivity of the performance to variations in resolution is examined. The performance of the schemes is both tested with idealized numerical tests and with a radon test for horizontal transport. For the latter test we have implemented the four new schemes in the global atmospheric tracer model TM2, described in [Hei95, VHK94]. We use one month of 12-hourly assimilated meteorological observations (ECMWF operational initialized analyses) at resolutions of $3.8° \times 5.0°$ and $2.5° \times 2.5°$ to drive the off-line model in order to compare the model results for different numerical schemes with measurements at two Indian Ocean sites. In order to investigate the sensitivity of model results to wind fields, time resolution of input meteorology, vertical resolution, and the parametrization of cumulus convection, we also perform runs with the TM3 model, an updated version of the TM2 model with respect to the vertical resolution and the cumulus parametrization, using 6-hourly analyses (ECMWF Re-Analysis) at a resolution of $3.8° \times 5.0°$. The islands used are

suitable for advection tests, because radon concentration peaks are mainly caused by transport and not by emission from the islands themselves, due to their small area.

This paper is organized as follows. The model structure and model equations are outlined in Section 6.2. In Section 6.3 we present the different numerical methods to solve the advection equation. The two test problems are described and discussed in Sections 6.4 and 6.5. We end with a conclusion in Section 6.6.

6.2 Model description and operator splitting

The basic equation in atmospheric transport modeling is the balance equation, which reads in pressure based, terrain-following hybrid coordinates,

$$\frac{\partial \rho\chi}{\partial t} + \frac{1}{a\cos\phi}\left[\frac{\partial(u\,\rho\chi)}{\partial\lambda} + \frac{\partial(v\,\cos\phi\,\rho\chi)}{\partial\phi}\right] + \frac{1}{h_\eta}\frac{\partial(w\,\rho\chi)}{\partial\eta} = R(\rho\chi) + S(\rho\chi). \quad (6.1)$$

The unknown χ denotes a vector of species mixing ratios, λ, ϕ and η are the coordinates in the longitudinal, latitudinal and vertical direction, u, v and w are the velocity components in the λ, ϕ and η direction, which are given, ρ is the density, a is the radius of the earth (6378 km), and h_η is a scale factor from the coordinate transformation. The term S represents the parametrized sub-grid scale transport in the vertical direction.

In this paper, the term R represents source and sink terms, but generally R represents a system of chemical reactions including source and sink. Thus in air quality models a lot of species are advected and stored.

6.2.1 Operator splitting

Equation (6.1) is normally solved using operator splitting, mainly to treat advection explicitly, and (stiff) chemistry implicitly. It also gives the possibility to use different time steps for different operators and different areas. We use the symmetric Strang [Str68] operator splitting, starting with an advection step of length $\frac{1}{2}\Delta t_{split}$, a step source and sinks terms and a convection step. Then we reverse the order and put two convection steps of $\frac{1}{2}\Delta t_{split}$ into one step of Δt_{split}. We use $\Delta t_{split} = 1$ hour.

6.2.2 Solving the sub-grid scale parametrization

The term $S(\rho\chi)$ in Equation (6.1) describes ECMWF's parametrizations for sub-grid scale vertical transport, based on the cumulus cloud calculation by the mass flux scheme of Tiedtke [Tie89] and the vertical diffusion coefficients calculated by the stability of the air using the formulae of Louis [Lou79]. The implementation of Tiedtke [Tie89] slightly differs between TM2 and TM3: in TM2 only updrafts are considered and in TM3 both updrafts and downdrafts are included. The sub-grid scale vertical transport parametrizations are the most important source of vertical

transport in the model. As in [Hei95] we integrate the part of Equation (6.1) that describes the sub-grid scale vertical transport over a grid box at height level k and obtain for each vertical column (i, j)

$$\frac{\mathrm{d}}{\mathrm{d}t} m_k \chi_k = \sum_{l=1}^{N_\eta} m_l \chi_l \left(f_{k-\frac{1}{2},l} - f_{k+\frac{1}{2},l} \right) \tag{6.2}$$

where m_k equals the air mass in layer k, $l = 1$ is the top level and $l = N_\eta$ is the ground level. The coefficients $f_{k+\frac{1}{2},l}$ represent the fraction of tracer mass of layer l that, per unit of time, crosses the layer boundary $k + \frac{1}{2}$ by means of sub-grid scale vertical transport processes. On the reduced grid these coefficients (available from meteorological input on the uniform grid) are averaged over the combined cells. We introduce $\vec{n}$, where the components $[\vec{n}]_k = m_k \chi_k$ are the tracer mass in a layer. We also introduce a matrix $\mathbf{M}$ with elements $[\mathbf{M}]_{k,l} = f_{k-\frac{1}{2},l} - f_{k+\frac{1}{2},l}$. Both have the dimension N_η. N_η represents the number of vertical layers, which equals 15 for the TM2 model and 19 for the TM3 model. So we have

$$\frac{\mathrm{d}}{\mathrm{d}t} \vec{n} = \mathbf{M} \cdot \vec{n}. \tag{6.3}$$

Equation (6.3) is solved fully implicitly

$$\vec{n}(t + \Delta t_S) = \vec{n}(t) + \Delta t_S \mathbf{M} \cdot \vec{n}(t + \Delta t_S) = (I - \Delta t_S \mathbf{M})^{-1} \vec{n}(t) = \mathbf{C} \cdot \vec{n}, \tag{6.4}$$

where Δt_S is equal to Δt_{split}. The matrix inversion, which gives the convection matrix $\mathbf{C}$, is done once per meteorological time step. The first-order moments are updated as described by Heimann [Hei95], and the second-order moments are treated analogously (see Section 6.3 for a short introduction and references on first and second-order moments). The assumption is that all moments involving horizontal components retain their information and are mixed as the means, and that the moments involving only vertical components gradually lose their information during convection [PMW+87].

6.2.3 Solving the source and sink terms

On the imposed grid, we have to solve in each grid point an Ordinary Differential Equation (ODE), which takes the specific form

$$m \frac{\mathrm{d}\chi}{\mathrm{d}t} = P - k_1 \, m\chi \tag{6.5}$$

with χ the mixing ratio of ^{222}Rn, P the emission (kg s^{-1}) and the rate constant $k_1 = 2.0974 \times 10^{-6}$ s^{-1}, which corresponds to the half-life time of 3.8 days. This ODE can be solved exactly:

$$m\,\chi(t + \Delta t_R) = \mathrm{e}^{-k_1 \Delta t_R} \, m\chi(t) + \left(1 - \mathrm{e}^{-k_1 \Delta t_R}\right) \frac{P}{k_1}, \tag{6.6}$$

with $\Delta t_R = \frac{1}{2}\Delta t_{split}$. The first and second-order moments are updated as follows: they all have the same decay factor $e^{-k_1 \Delta t_R}$, and the first-order moment in the vertical direction at ground level is adjusted to cope with emission at the bottom.

6.3 Numerical methods: advection schemes

We are comparing six numerical schemes for advection: two schemes presently used in the TM model, namely Slopes developed by Russell and Lerner [RL81] and Second Moment developed by Prather [Pra86], and four recently developed schemes: Mol-rg, Split-u, Split-us and Split-rg. These new schemes use third-order calculations of fluxes at cell interfaces, together with flux-limiting, based on formulas of van Leer [vL77], and Koren [Kor93]. These limiters are necessary to prevent over- and undershoot. The 2D versions of these new schemes have been tested in [HS95, BHV94, Spe95]. To extend these schemes to 3D, fluxes based on concentrations are replaced by mixing ratio fluxes. A precise mathematical description of these mixing ratio fluxes is given in Appendix 6.7.1. For ease of comparison we also include in the tests the Donor Cell scheme on a reduced grid.

We focus on mass conservative schemes. The above schemes are all formulated in terms of in- and outgoing fluxes, and are therefore strictly mass conservative, in contrast with e.g. semi-Lagrangian schemes. Semi-Lagrangian schemes are therefore not included in the tests. An indirect comparison of Split-us with the semi-Lagrangian TREMBA schemes of Smolarkiewicz and Rasch [SR91] can be found in [HS95]. This comparison with limited TREMBA schemes was favorable for Split-us, especially with respect to cpu time and mass errors. With semi-Lagrangian schemes limiting is also necessary to achieve positivity.

The main difference between the advection schemes is the interpolation of the mixing ratio to obtain the fluxes at cell interfaces from cell center values. Donor Cell and the new schemes use only cell center values. Donor Cell is a first-order upwind scheme, and the new schemes are third-order upwind schemes with flux-limiting. To compute the fluxes, Slopes uses also the first-order moment (slope) of the mixing ratio, and Second Moment uses the first and second-order moments of the mixing ratio (which explains their names). This results in different memory requirements for the last two schemes, see Section 6.4.

Slopes, Second Moment, Split-u, Split-us and Split-rg are so-called dimensional splitting schemes. In these schemes one 3D advection step is a sequence of 1D advection steps. The sequence is chosen such that we have a symmetric Strang [Str68] splitting and for each direction a maximum Courant Number (defined in Appendix 6.7.1) that is as close as possible to, but smaller than 1. Slopes, Second Moment and Split-u divide one advection step in four longitudinal steps, two latitudinal steps and one step in the vertical direction. Split-rg and Split-us divide one advection step in two steps for both the longitudinal and latitudinal direction, and one step in the vertical direction. Close to the polar regions, the advection step in the longitudinal direction

is further subdivided for Slopes and Second Moment.

We have two Method of Lines (MOL) schemes in our test: The Donor Cell scheme and Mol-rg. In the MOL approach the space and time are discretized separately. The space discretization leads to an ODE which is integrated with a Runge-Kutta scheme. In this approach all three space dimensions are kept together. This results in one single 3D advection time step for the whole domain, including the polar regions. In Mol-rg we have taken a standard second-order Runge-Kutta scheme for which monotonic results are guaranteed for Courant numbers up to $\frac{1}{2}$; in practice this Courant restriction can be relaxed to $\frac{2}{3}$, see [HKvLV95]. On a uniform grid, the Courant restriction imposes a very small time step. In the MOL approach we therefore use a reduced grid.

Three schemes use a reduced grid: Mol-rg, Split-rg and Donor Cell. The reduced grid [PMW+87, Wil92, Ras94] is a uniform grid with less cells near the poles to overcome the polar singularity, see Figure 6.1 for an example. Grid reduction means that at a small number of latitudes near the poles the grid size in the longitudinal direction is doubled.

The characteristics of the new schemes are:

Mol-rg A Method of Lines (MOL) scheme on a reduced grid. Mol-rg is discretized with third-order upwind fluxes on a reduced grid, and uses a second-order Runge-Kutta scheme viz. the explicit trapezoidal rule, for time integration [HKvLV95].

Split-u a dimensional splitting scheme on a uniform grid, based on formulas derived in [HS95]. The 1D advection steps are solved in conservative form with third-order upwind biased fluxes together with flux-limiting using direct discretization [HT94]. Due to the term $\cos\phi$ in Equation 6.1, the CFL condition requires very small time steps on fine grids which makes this scheme very expensive in terms of cpu time. Therefore two modifications of Split-u are developed.

Split-us a modification of Split-u, where the longitudinal direction is made unconditionally stable by allowing the stencil to vary with the courant number [HS95]. This is as with semi-Lagrangian methods, but Split-us is formulated in terms of in- and outgoing fluxes and therefore we maintain the mass-conservation property.

Split-rg a modification of Split-u, where the uniform grid is replaced by a reduced grid. Due to the reduced grid, relatively large time steps can be used. As contrasted with Split-us, this scheme has good efficiency properties on a vector computer.

The schemes Mol-rg, Split-us and Split-rg are available through WWW or by contacting the first author.

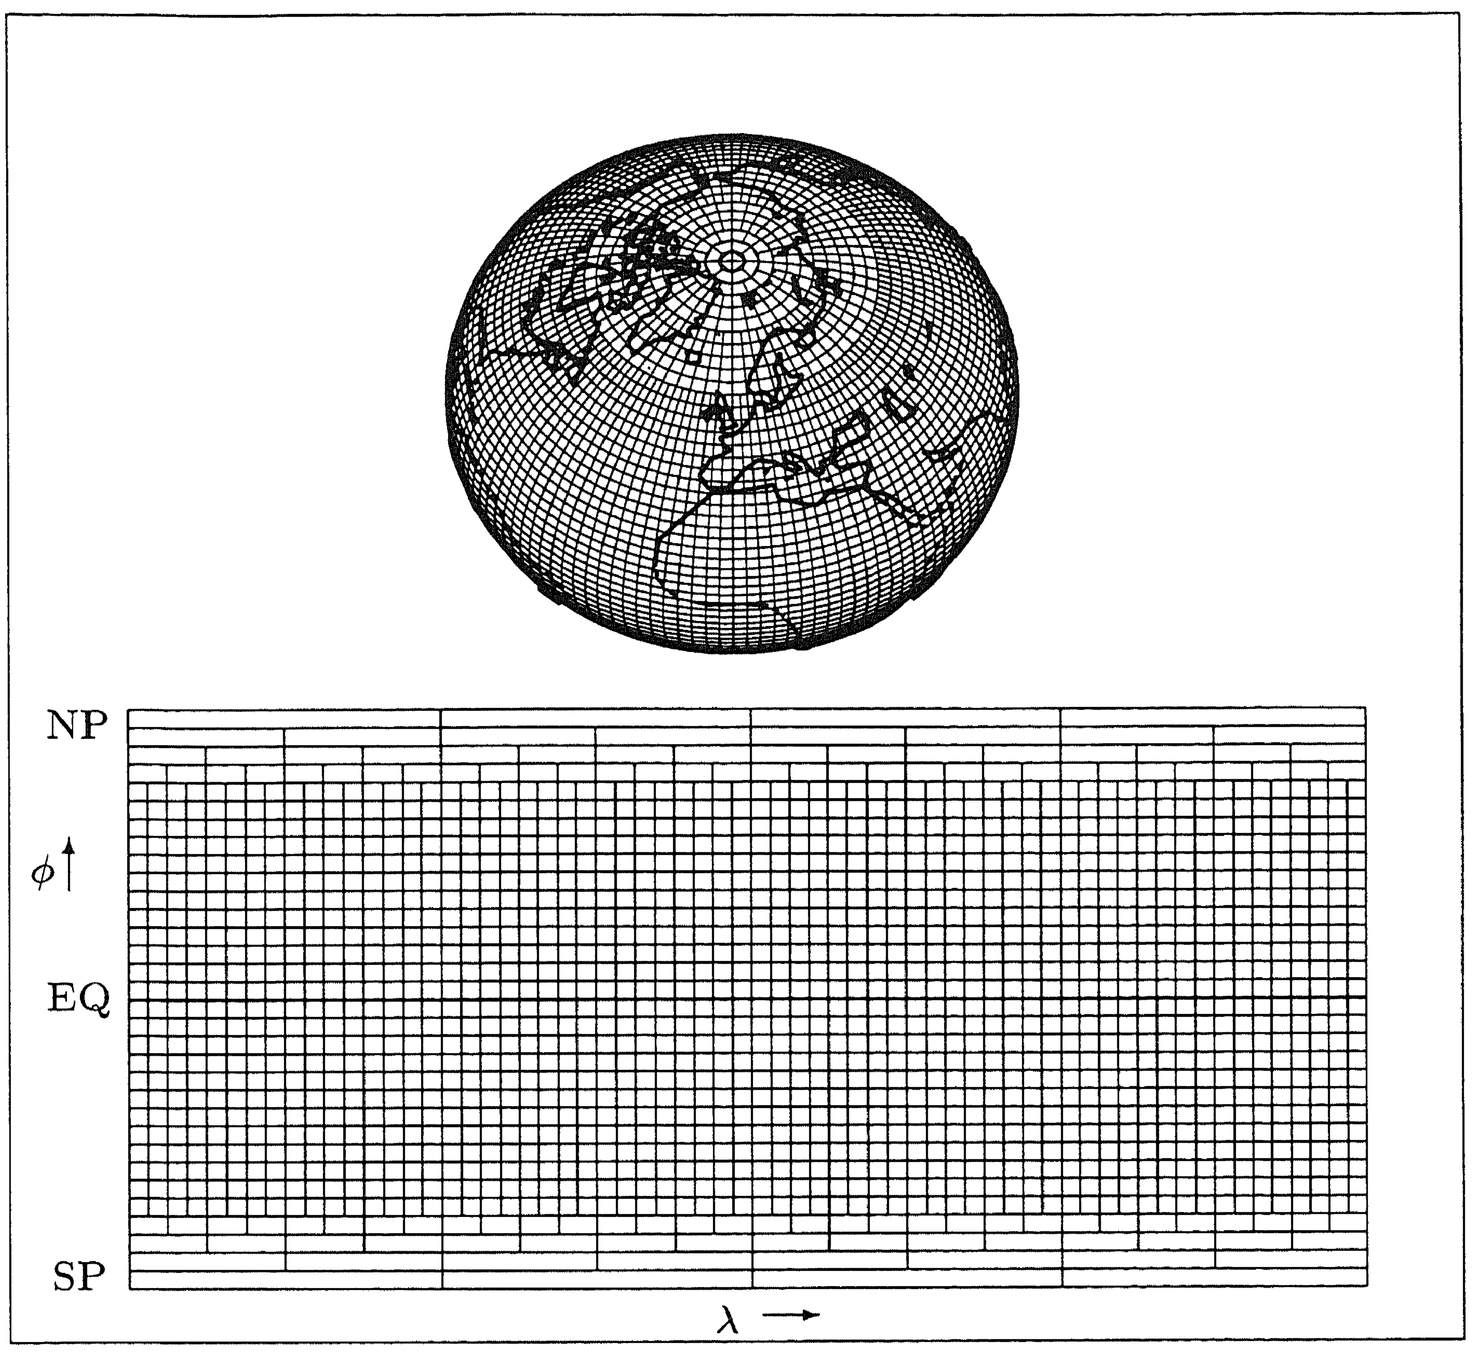

Figure 6.1: A reduced grid on the globe (top), and a 64 × 32 reduced grid in longitude/latitude coordinates (bottom).

6.4 Problem I: Solid-body rotation on a sphere

6.4.1 Introduction

The solid-body rotation is a well known test problem for advection schemes [WR89]. In this test we only have horizontal advection, so we solve the model equation (6.1) without the terms S and R and with the vertical velocity component equal to zero. As in [SR91, HS95], we consider a cone-shaped initial profile c_0, a cylinder-shaped initial profile c_1 and a smooth initial profile c_2 given by

$$c_0(\lambda,\phi) = \max(0, 1 - r(\lambda,\phi)/R), \tag{6.7a}$$

$$c_1(\lambda,\phi) = \begin{cases} 1 & \text{if } r(\lambda,\phi) > R, \\ 2 & \text{if } r(\lambda,\phi) \le R, \end{cases} \tag{6.7b}$$

$$c_2(\lambda,\phi) = \cos^4\left(\lambda - \tfrac{1}{2}\pi\right)\cos^4(\phi), \tag{6.7c}$$

where

$$r(\lambda,\phi) = 2\sqrt{\cos^2(\phi)\sin^2\left(\tfrac{1}{2}\left(\lambda - \tfrac{3}{2}\pi\right)\right) + \sin^2\left(\tfrac{1}{2}\phi\right)},$$

and $R = 7\pi/64$. With the smooth profile we can test the convergence behavior, the cylinder profile with background is useful to test whether or not the advection schemes are monotonic, and the cone profile shows the amount of smearing.

The velocities are given by

$$u = 2\pi(\cos\beta\cos\phi + \sin\beta\sin\phi\cos\lambda), \quad v = -2\pi\sin\beta\sin\lambda, \tag{6.8}$$

where $\beta = 90°$. Observe that the wind velocities are constant in time. With this wind field, the exact solution after a full rotation is equal to the initial profile. This wind field is numerically divergence free, if the grid sizes in the λ and ϕ direction are equal.

6.4.2 Results for problem I

The results for the solid-body rotation test are given in Tables 6.1, 6.2, and 6.3 for the respective resolutions 5.6°× 5.6°, 2.8°× 2.8°, and 1.4°× 1.4°. We use the same errors as Smolarkiewicz and Rasch [SR91], defined by

$$emin = \frac{\min(c^n_{i,j}) - \min(c_{i,j}(t_n))}{\max(c_{i,j}(t_n))}, \qquad emax = \frac{\max(c^n_{i,j}) - \max(c_{i,j}(t_n))}{\max(c_{i,j}(t_n))},$$

$$err0 = \frac{(\sum \gamma_j (c^n_{i,j} - c_{i,j}(t_n))^2)^{1/2}}{\max(c_{i,j}(t_n))},$$

where $\gamma_j = \cos(\phi_j)/(N_\lambda \sum_{k=1}^{N_\phi} \cos(\phi_k))$. In all formulas $c_{i,j}(t_n)$ denotes the reference solution, and $c^n_{i,j}$ the computed solution at time t_n, the time necessary for a full

scheme	cpu	nt	cpu'	undershoot	overshoot	smearing	$L2$
	s		ms				
Second M.	1.8	82	10.5	-2.26E-2	1.50E-2	-1.71E-1	1.55E-2
Slopes	0.7	82	4.0	-1.89E-2	4.16E-2	-2.57E-1	1.42E-2
Split-u	1.8	163	5.2	-1.45E-13	-3.29E-2	-4.72E-1	8.61E-3
Split-us	0.4	48	4.0	-2.71E-13	-2.59E-3	-3.77E-1	6.22E-3
Split-rg	0.7	78	4.5	4.32E-14	-9.09E-2	-5.73E-1	1.67E-2
Mol-rg	1.2	263	1.1	-2.82E-14	-1.42E-1	-6.48E-1	2.07E-2
Donor Cell	0.2	174	0.4	1.90E-7	-3.77E-1	-9.00E-1	1.48E-1

Table 6.1: Results for solid-body rotation test on a sphere for the coarse grid ($\approx$ 5.6° $\times$ 5.6°): Efficiency in terms of cpu time versus accuracy in terms of undershoot and overshoot for a cylinder with background, smearing for cone and $L2$ for a smooth profile. Small typed numbers are close to round-off. cpu' equals cpu / $(nt\, N_\lambda\, N_\phi)$

scheme	cpu	nt	cpu'	undershoot	overshoot	smearing	$L2$
	s		ms	$emin\ c_1$	$emax\ c_1$	$emax\ c_0$	$err0\ c_2$
Second M.	37.7	326	13.9	-3.36E-2	3.65E-2	-9.05E-2	5.45E-3
Slopes	10.5	326	3.9	-2.14E-2	3.89E-2	-1.37E-1	4.86E-3
Split-u	29.4	652	5.4	-1.15E-12	-2.26E-6	-2.22E-1	1.61E-3
Split-us	3.1	96	3.9	-6.71E-13	-2.56E-9	-1.66E-1	1.35E-3
Split-rg	3.8	108	4.2	-1.44E-13	-8.65E-5	-2.82E-1	3.16E-3
Mol-rg	6.8	415	2.0	-1.36E-16	-6.50E-3	-3.75E-1	4.06E-3
Donor Cell	1.2	274	0.5	2.00E-15	-2.72E-1	-8.13E-1	8.53E-2

Table 6.2: As Table 6.1, but now for the middle grid ($\approx$ 2.8° $\times$ 2.8°).

scheme	cpu	nt	cpu'	undershoot	overshoot	smearing	$L2$
	s		ms	$emin\ c_1$	$emax\ c_1$	$emax\ c_0$	$err0\ c_2$
Second M.	3039	1304	70.6	-3.17E-2	3.63E-2	-4.61E-2	1.91E-3
Slopes	351	1304	8.2	-2.93E-2	3.54E-2	-7.90E-2	1.69E-3
Split-u	671	2608	7.8	-4.64E-12	1.64E-14	-1.17E-1	3.38E-4
Split-us	36	192	5.7	-2.77E-12	7.31E-14	-9.29E-2	3.31E-4
Split-rg	60	325	5.6	-6.49E-13	-6.88E-11	-1.50E-1	6.52E-4
Mol-rg	115	993	3.5	-2.10E-13	-1.07E-8	-1.65E-1	7.32E-4
Donor Cell	31	724	1.3	-6.84E-14	-1.65E-1	-6.97E-1	5.51E-2

Table 6.3: As Table 6.1, but now for the finest grid ($\approx$ 1.4° $\times$ 1.4°).

rotation. The number of steps necessary to fulfill the CFL condition for a full rotation is nt.

The error *emin* equals the lowest concentration in the numerical solution minus the lowest concentration in the reference solution. The error *emax* is the same difference for the highest concentration. Because advection does not generate new extrema, *emin* should be less or equal to zero, and *emax* should be greater then or equal to zero. The error $err0$ is a root-mean-square error. The scaling is chosen such that the errors will be equal to 1 if the error is 1 in all grid points and $\max(c_{i,j}(t_n)) = 1$.

The smearing is given by the *emax* with the cone c_0 as initial profile and $L2$ is the $err0$ error for the smooth initial profile c_2. We define overshoot as the *emax* error and undershoot as the *emin* error, both with the cylinder c_1 as initial profile.

The *emax* error with the cylinder c_1 is influenced by overshoot and numerical diffusion. The influence of numerical diffusion decreases with increasing resolution. In this test we see that for the new schemes this error is purely caused by numerical diffusion, and not by overshoot, so we conclude that these schemes are monotonic. The limiter used in Slopes and Second Moment prevents only negative values, resulting in 3% overshoot and undershoot.

The Split-us scheme gives very good results for this test. This may not be the case in a full-scale chemistry transport model, because the time step can be determined by other processes, so that Split-us must take smaller time steps than necessary from a point of stability. The timing is done on a workstation and we know that on a vector processor such as the Cray C90, the efficiency for the flux calculation in the longitudinal direction is low. This is caused by the summation over c_l in Equations (3.12) and (3.13) in Chapter 3 of this thesis. The difference between $L2$ for uniform and reduced grid is quite large, which is caused by the fact that in this test the center of the transported concentrations is relatively long in the polar regions. If in actual computations the polar regions are of special interest, one could start grid reduction at e.g. 70° instead of 60°. Note that also Split-rg can be made unconditionally stable, with as disadvantage a low efficiency on a vector processor.

Differences in cpu time are mainly caused by nt. For Split-us and the schemes on a reduced grid a doubling of the resolution leads to a doubling of nt, whereas on a uniform grid nt quadruples for explicit finite volume schemes, due to the extra factor 2 coming from the $\cos\phi$. So on high resolutions, a reduced grid or unconditional stability is necessary to avoid very small time steps. We prefer the reduced grid approach, because all other processes need about 25% less cpu time due to the lower number of cells. Other differences in cpu time are caused by the number of flux calculations per step (7 for Slopes, Second Moment and Split-u, 6 for Mol-rg, 5 for Split-us and Split-rg, and 3 for Donor Cell), and the costs of one flux calculation, which is relatively high for limited schemes. Timings also depend on implementation, of course.

Slopes and Second Moment need respectively 4 to 10 times more memory, due to the storage of the first and second-order moments of the transported species. This will

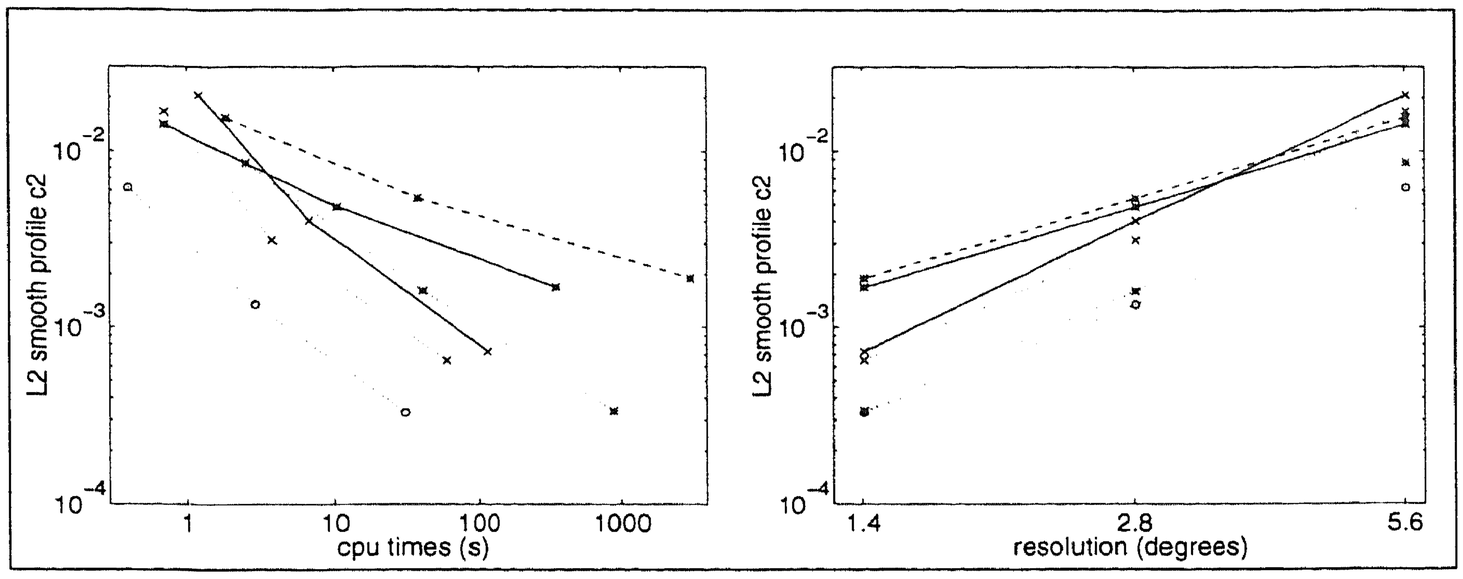

Figure 6.2: Convergence behavior. Bottom: $L2$ as a function of cpu time. Top: $L2$ as a function of the resolution. The new schemes have a higher convergence rate than Slopes and Second Moments. The convergence rate is the increment in accuracy when the resolution becomes finer. *-solid: Slopes, *-dashed: Second Moment, *-dotted: Split-u, o-dotted: Split-us, x-dotted: Split-rg, x-solid: Mol-rg. Results for Donor Cell fall outside these plots.

hardly change when using Second Moment on a reduced grid, as in [PMW+87]. In our hardware configuration, the moments gives problems with the data-cache, which is the reason why the ratio cpu / $(nt\,N_\lambda\,N_\phi)$ is not constant for a given method when increasing the resolution. In realistic models with many chemical species this memory requirement will be a major draw-back.

We observe that Split-us, Split-rg and Mol-rg give good results for smooth profiles and therefore good convergence speed, see Figure 6.2. Also with block profiles the new schemes give good results. Slopes and Second Moment give good results for the cone tests.

We conclude that Split-u and Second Moment are very expensive in terms of cpu time, and that Second Moment is also very expensive in terms of memory requirements. Results for Split-u will not be given in the radon test, but results for Second Moment are given for reasons of comparison.

6.5 Problem II: Radon test

6.5.1 Introduction

We use the radioactive noble gas radon-222 (^{222}Rn) to perform a horizontal transport experiment that allows us to compare the different advection schemes under real atmospheric conditions. We assume that ^{222}Rn is emitted with a constant flux of 1 atom cm^{-2} s^{-1} from the land surface and no flux from the ocean. Heimann et al.

[HMP90] and Balkanski and Jacob [BJ90] observe that the value of 1 atom cm^{-2} s^{-1} has an uncertainty of approximately 25%. Since the flux from the ocean is two to three orders of magnitude smaller, we can neglect the contribution from the ocean. When the land is covered with a permafrost layer or snow the emission flux is zero. For January we use a climatological value of 60°N for the latitude circle that bounds the area covered with permafrost or snow layers. The half-life of ^{222}Rn is 3.8 days.

Two datasets of ECMWF operational initialized analyses preprocessed by the KNMI for the month January 1992 are used as input to the 3D chemical transport model. Pressure level fields from ECMWF's operational archive with a vertical resolution of 15 layers and transformed to a uniform 2.5° × 2.5° grid were used as input to the preprocessing procedure. The time resolution of both preprocessed datasets is 12 hours, the spatial resolution is respectively 3.8° × 5.0° and 2.5° × 2.5° in the horizontal and 15 layers in the vertical direction [VHK94]. In order to investigate the sensitivity to the version of the ECMWF forecast model and to the vertical resolution of the preprocessed data, we also use a third dataset of the ECMWF Re-Analysis (ERA) preprocessed by the KNMI for the same month. In the ERA project the ECMWF forecast model version of April 1995 was used to re-analyze the 15-year period 1979-1993 [GÅU96]. The differences between the operational analysis and re-analysis datasets are due both to the inclusion of more observations (e.g. weekly global sea surface temperature analyses) and changes in physical parametrizations over the period from 1992 to 1995 (e.g. introduction of prognostic equations for cloud water content and cloud cover). The time resolution of the preprocessed re-analysis dataset is 6 hours and the spatial resolution is 3.8° × 5.0° with 19 layers in the vertical direction. All datasets contain the grid-resolved mass fluxes and the input necessary to calculate sub-grid scale vertical transport. Similar preprocessing routines as described for the standard version of the TM2 model by Heimann [Hei95] were used in constructing the datasets.

In this radon test the differences between the advection schemes become apparent in the maximum surface concentrations, 'plume-around-maximum' areas, surface concentration front patterns near coastlines, and the structure of 'radon concentration peaks' associated with rapid advection of boundary-layer air from continents to remote oceanic sites. Numerical diffusion and numerical dispersion are critical concerns in the simulation of oceanic radon concentration patterns [BC89, BJ90]. In the world's oceans there are a few remote island sites where the surface ^{222}Rn concentration is continuously measured. We have selected only two small islands, namely Amsterdam Island (77°34'E, 37°50'S) and Crozet Island (51°52'E, 46°26'S) in the southern Indian Ocean, where the influence of local radon emissions has been shown to be rather small [HMP90]. We use the data of these two islands for January 1992 in order to compare the different numerical simulations with actual measurements.

6.5.2 Results for the radon test

From the comparison between measurements and numerical model calculations,

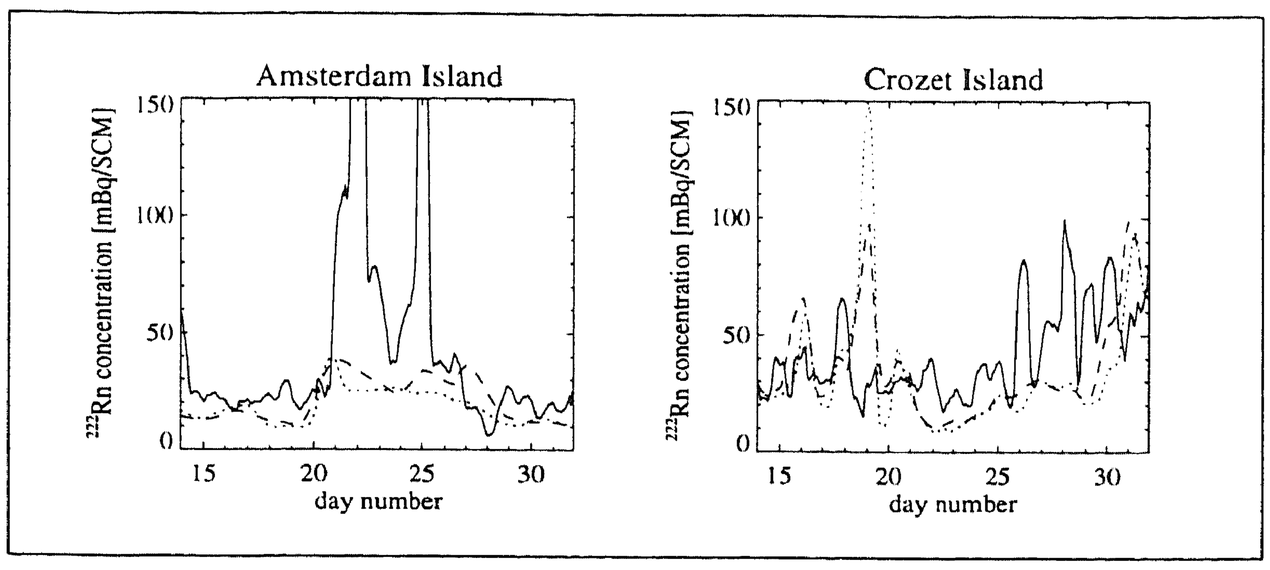

Figure 6.3: Smoothed local radon concentration measurements (solid line) and numerical model results at two resolutions, respectively 2.5° × 2.5° (dotted) and 3.8° × 5.0° (dashed), from 0000 UTC 14 January 1992 to 2400 UTC 31 January 1992. The numerical scheme in this plot is the Split-rg scheme. The lowest model grid cell containing the measurement points is used. One Standard Cubic Meter (SCM) of air is defined as one cubic meter of air at standard temperature (273 K) and standard pressure (1013 mbar).

plotted in Figure 6.3, we can see that for both resolutions the model seems to show only a small systematic bias in the predicted concentrations. The mean model concentrations are somewhat lower than the measurements. Heimann et al. [HMP90] found too high model concentrations at all stations using ECMWF wind fields at a resolution of 7.8° × 10.0° and the Slopes scheme, while Balkanski and Jacob [BJ90] using GISS GCM wind fields at a resolution of 3.8° × 5.0° and the Second Moment scheme did not simulate such an overprediction. Note that we compare the model results with the measurements only for the last 18 days, because the model results for the first two weeks strongly depend on the initial concentration used.

From Figure 6.3 it can be concluded that only crude features in the measurements are captured by the model results at both model resolutions, namely the higher concentration at Amsterdam Island in the middle of the 18-day period and the higher concentration at Crozet Island at the end of the month. At the 3.8° × 5.0° model resolution the correlation between the measurements (12-hourly smoothed) and the model results is significant for both islands, but at the 2.5° × 2.5° resolution the correlation is significant only for Amsterdam Island. In order to determine the significance we have used 36 12-hourly samples in the last 18 days of the month and adopted a 95% confidence level. The correlation coefficients are 0.45 and 0.27 for respectively Amsterdam and Crozet Islands at the low resolution, and 0.35 and 0.08 at the high resolution. The fact that the cross-correlation is worse for finer resolution only indicates that errors in the input wind fields and sub-grid scale transport become more apparent for these finer resolutions. As we will discuss below, the structure of the

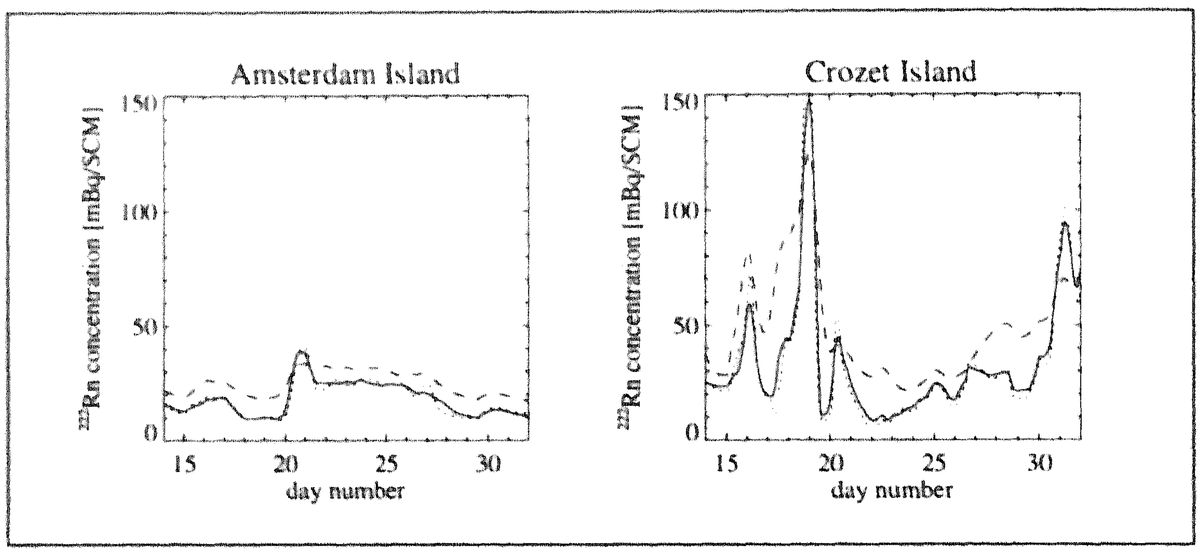

Figure 6.4: Numerical model results at a resolution of 2.5° × 2.5° for different numerical schemes from 0000 UTC 14 January 1992 to 2400 UTC 31 January 1992. Results are plotted for the Split-rg (solid line), Slopes (dotted), and Donor Cell (dashed) schemes. The lowest model grid cell containing the measurement points is used.

radon concentration peaks is better in the 2.5°× 2.5° resolution version of the model than in the 3.8°× 5.0° resolution version.

Although the correlation for three of the four runs is significant, the model prediction of the height and timing of radon concentration peaks is poor for this 18-day period. For Amsterdam Island the measured radon concentration peaks on 21/22 January and 24/25 January do not appear similarly in the model. For Crozet Island no radon concentration peaks were measured while the higher-resolution model simulates a relatively high radon concentration peak on 18 January.

In Figures 6.4 and 6.5 model results at the two resolutions are plotted. For both stations and both resolutions the strong diffusivity of the Donor Cell scheme causes significantly higher background concentrations over the oceans than as simulated with the other, less diffusive, numerical schemes. However, even the Donor Cell scheme does not systematically overpredict the concentrations. The similar background concentrations simulated by the Slopes scheme and the Second Moment scheme (the last one is not shown) confirm the suggestion by Heimann et al. [HMP90] that the difference between their model results and those of Balkanski and Jacob [BJ90] are mainly due to differences in resolution and not in numerical schemes. The results for the Split-us scheme coincide with the results for the Split-rg scheme, and are therefore not shown.

To evaluate the general capability of the model to simulate radon concentration peaks (looking now at their structure and not their timing) we compared autocorrelation functions for both model resolutions with the autocorrelation function for the 12-hourly smoothed measurements (not shown). Although an 18-day period is too short to formulate firm conclusions about radon concentration peak statistics, we find

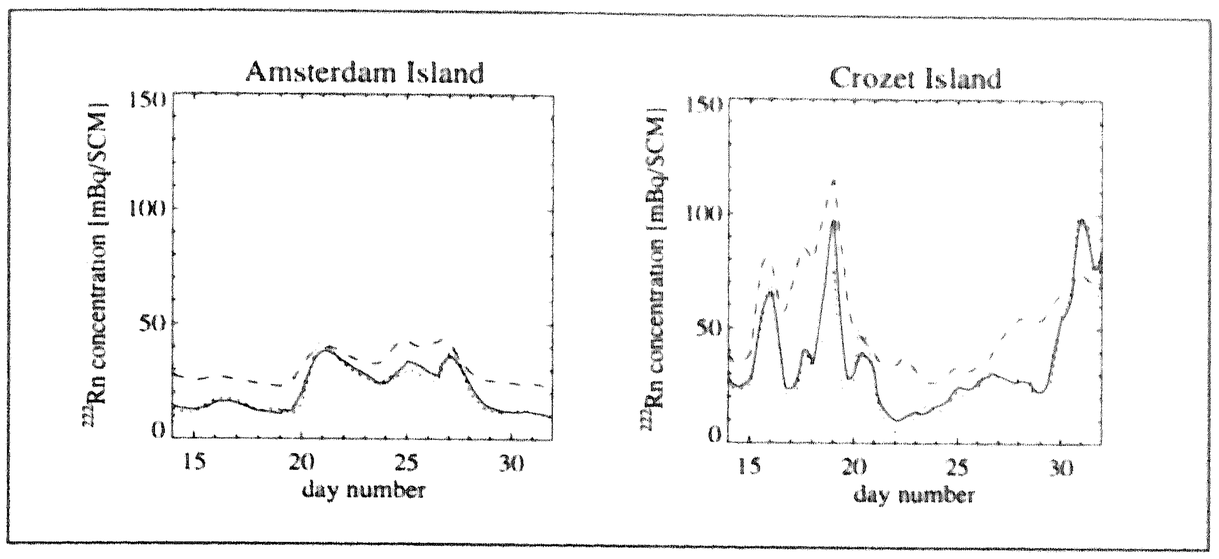

Figure 6.5: Numerical model results at a resolution of 3.8° × 5.0° for different numerical schemes. Same plot as Figure 6.4.

that for the 2.5° × 2.5° resolution the autocorrelation functions generally fall off to 0.5 at a similar time lag of approximately 15 hours as for the smoothed data (here excluding the Donor Cell scheme). In [HMP90], using a whole year in the calculation of the autocorrelation, the time lag at an autocorrelation of 0.5 was found to be twice as large in their model as in their smoothed data. We find the same difference between our 3.8° × 5.0° resolution results and the data, indicating that a spatial resolution of 2.5° × 2.5° is needed to adequately represent time scales of the order of 12 hours in atmospheric transport models. This conclusion is related to the flow structure in the lower troposphere and is independent of numerics. We want to emphasize here that Jacob et al. [Jc97] only refer to the frequency and magnitude of radon concentration peaks, but not to their structure, when they conclude that the synoptic-scale transport of continental plumes is resolved with no significant numerical diffusion. Here we find that numerical diffusion also has no significant influence on the structure of radon concentration peaks, which are resolved provided that a resolution of 2.5°× 2.5° is used.

Coming back to the difficulty the model has to correctly predict the height and timing of the radon concentration peaks, we can identify three causes for the transport model errors. One is the error in the lower tropospheric circulation pattern as analyzed by ECMWF. This is due to a lack of meteorological measurements in the southern Indian Ocean region that can be used in the data assimilation procedure at ECMWF. Secondly, the results can be sensitive to the parametrization of sub-grid scale vertical transport. A final source of errors is the use of a 12-hour time resolution of the input meteorology instead of using a higher time resolution, for example by doing on-line tracer transport in a climate model nudged towards observed meteorology.

In order to assess the three mentioned sources of error for this particular experiment, we have performed additional runs with the TM3 model. From Figure 6.6

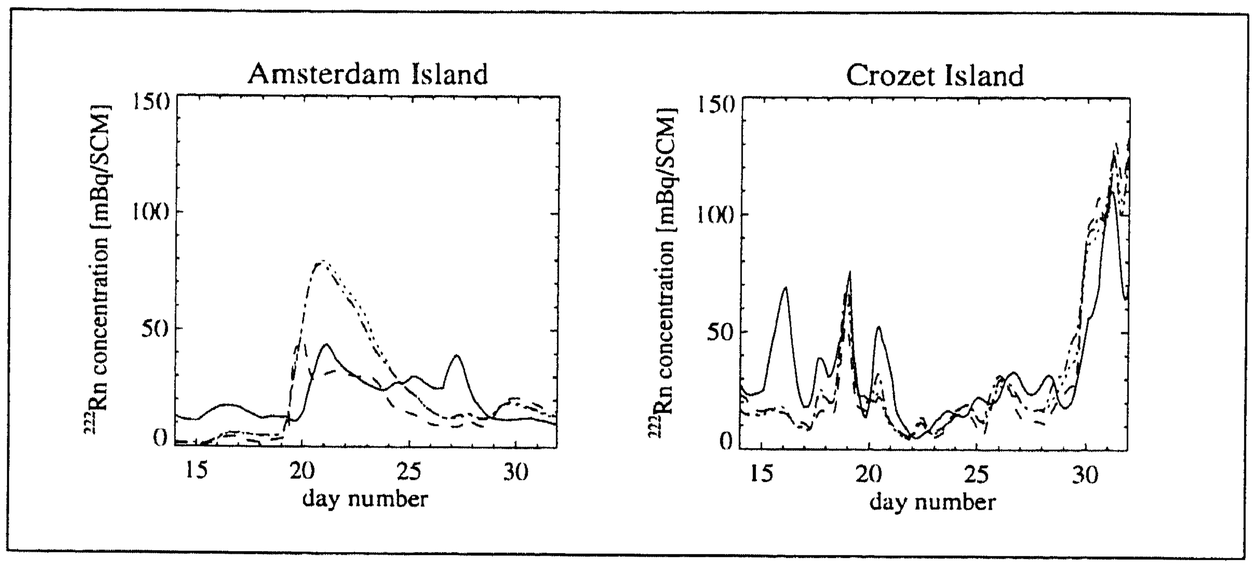

Figure 6.6: Numerical model results at a resolution of 3.8° × 5.0° for the sensitivity tests. Shown are the results of: the TM2 model (solid line; identical to the dotted line in Figure 6.5), the standard run with the TM3 model (dotted), TM3 with cumulus convection turned off (dashed), and TM3 with a reduction of the number of layers close to the surface—bottom 4 layers reduced to 2 layers—(dash-dot). All results are obtained with Slopes. The lowest model grid cell containing the measurement points is used.

we can conclude that the results significantly differ for the TM2 and TM3 models. Probably differences in the sub-grid scale vertical transport are responsible for the difference between the models, since the wind fields in this region are found to be almost identical for the operational (1992) and the re-analysis ECMWF models (see Figure 6.7). We did two sensitivity experiments (also shown in Figure 6.6): one in which the cumulus convection was switched off and one in which the boundary-layer resolution was reduced. Our first sensitivity study reveals that cumulus convection can have an effect on near-surface radon concentrations over the oceans (cf. Brost and Chatfield [BC89]). We furthermore show in Figure 6.6 that reducing the vertical resolution in the boundary layer by approximately a factor of two, and keeping the same vertical mass exchange as with the higher vertical resolution, does not have any influence on the model results at the two islands. Summarizing we can say that the TM3 model results agree slightly better with the measurements, but still the height and timing of radon concentration peaks is not correctly predicted. So, the fact that we have a 6-hour time resolution of input meteorology in TM3 does not lead to a significant improvement of the modeling performance at these islands.

From Figure 6.7 we conclude that during the whole month there are large differences between the measured local wind directions at the islands and the wind directions analyzed by the ECMWF in the lowest grid cells containing the respective islands. Although the ECMWF wind field in a grid cell does not capture local effects on the small islands, the fact that we observe persistent deviations for different wind

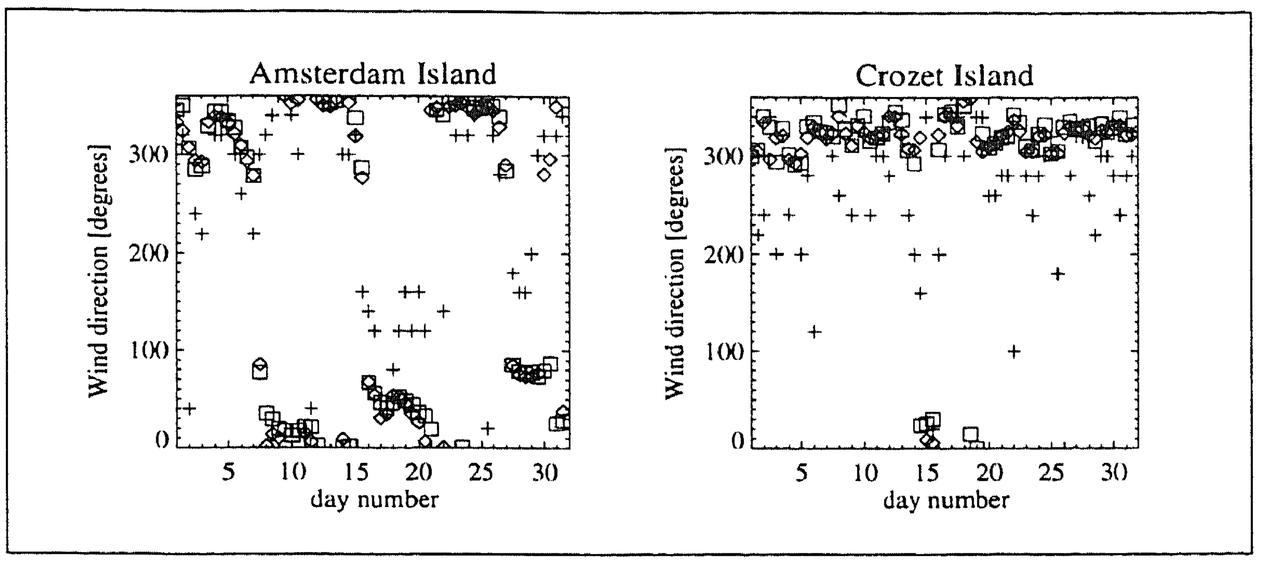

Figure 6.7: Measured wind directions at the surface on the islands (crosses) and near-surface wind directions calculated from ECMWF operational initialized analyses (diamonds) and the ECMWF Re-Analysis (squares), all plotted at 12-hour time resolution. The lowest model grid cell containing the measurement points is used. A wind direction of zero degrees corresponds to the north-to-south direction.

directions allows us to conclude that the analyzed wind directions in this Indian Ocean region contain considerable errors. We therefore think that using a higher time resolution in a climate model nudged towards the ECMWF analyses, would likely have produced a similar erroneous result. However, from this limited period no general conclusions regarding the respective contributions of the three error sources can be drawn, and we must stress that this radon test as a whole is not conclusive with respect to the error sources.

Figure 6.8 gives the average ^{222}Rn concentration for the last 18 days of January 1992 in a part of the Indian Ocean and Southern Africa for the advection schemes Donor Cell, Split-rg and Slopes at both resolutions. Figure 6.9 gives the instantaneous ^{222}Rn concentration at 0000 UTC 21 January 1992. From both figures it can be concluded that Split-rg and Slopes are much less diffusive than the Donor Cell scheme, and that the model results for Split-rg and Slopes are almost identical. This is also the case for the Mol-rg, Split-us and Second Moment schemes (not shown). Still, focusing now on the results for Split-rg and Slopes shown in Figure 6.9, we find small differences in the instantaneous surface concentration front patterns and the associated shape of radon plumes advected towards the Indian Ocean islands. We also see that Slopes has somewhat higher maxima. For the average concentrations shown in Figure 6.8 the differences are even much smaller over the ocean and somewhat smaller over the continent, compared to the instantaneous results shown in Figure 6.9. In order to quantify the differences between the averaged results for Split-rg and Slopes, we have determined the maximum value of the average ^{222}Rn concentration over the whole globe. Slopes gives 11.8 and 11.6 Bq/SCM for respectively the fine and coarse resolution, where Split-rg gives 10.2 and 10.6 Bq/SCM. This maximum is found in

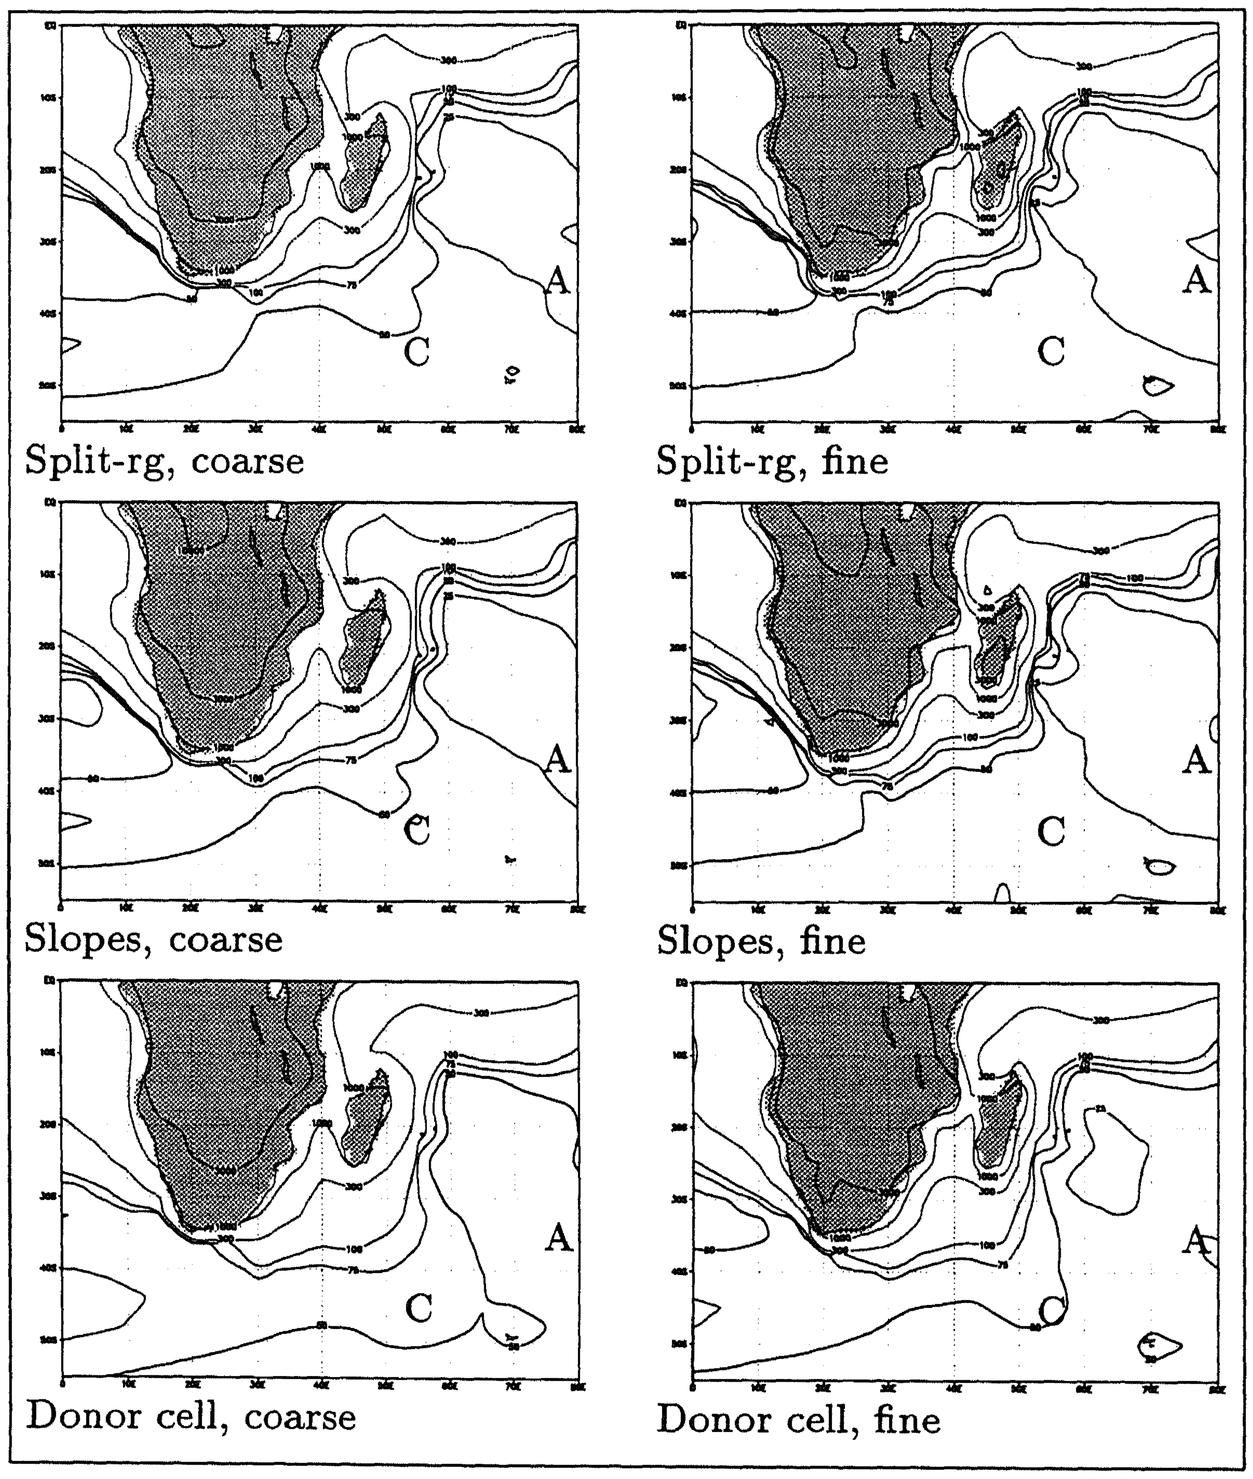

Figure 6.8: Contourlines of 18-day average ^{222}Rn concentration in mBq/SCM in a part of the Indian Ocean and southern Africa, for three advection schemes and two resolutions (coarse: $3.8° \times 5.0°$, fine: $2.5° \times 2.5°$). A and C indicate the location of Amsterdam Island and Crozet Island respectively.

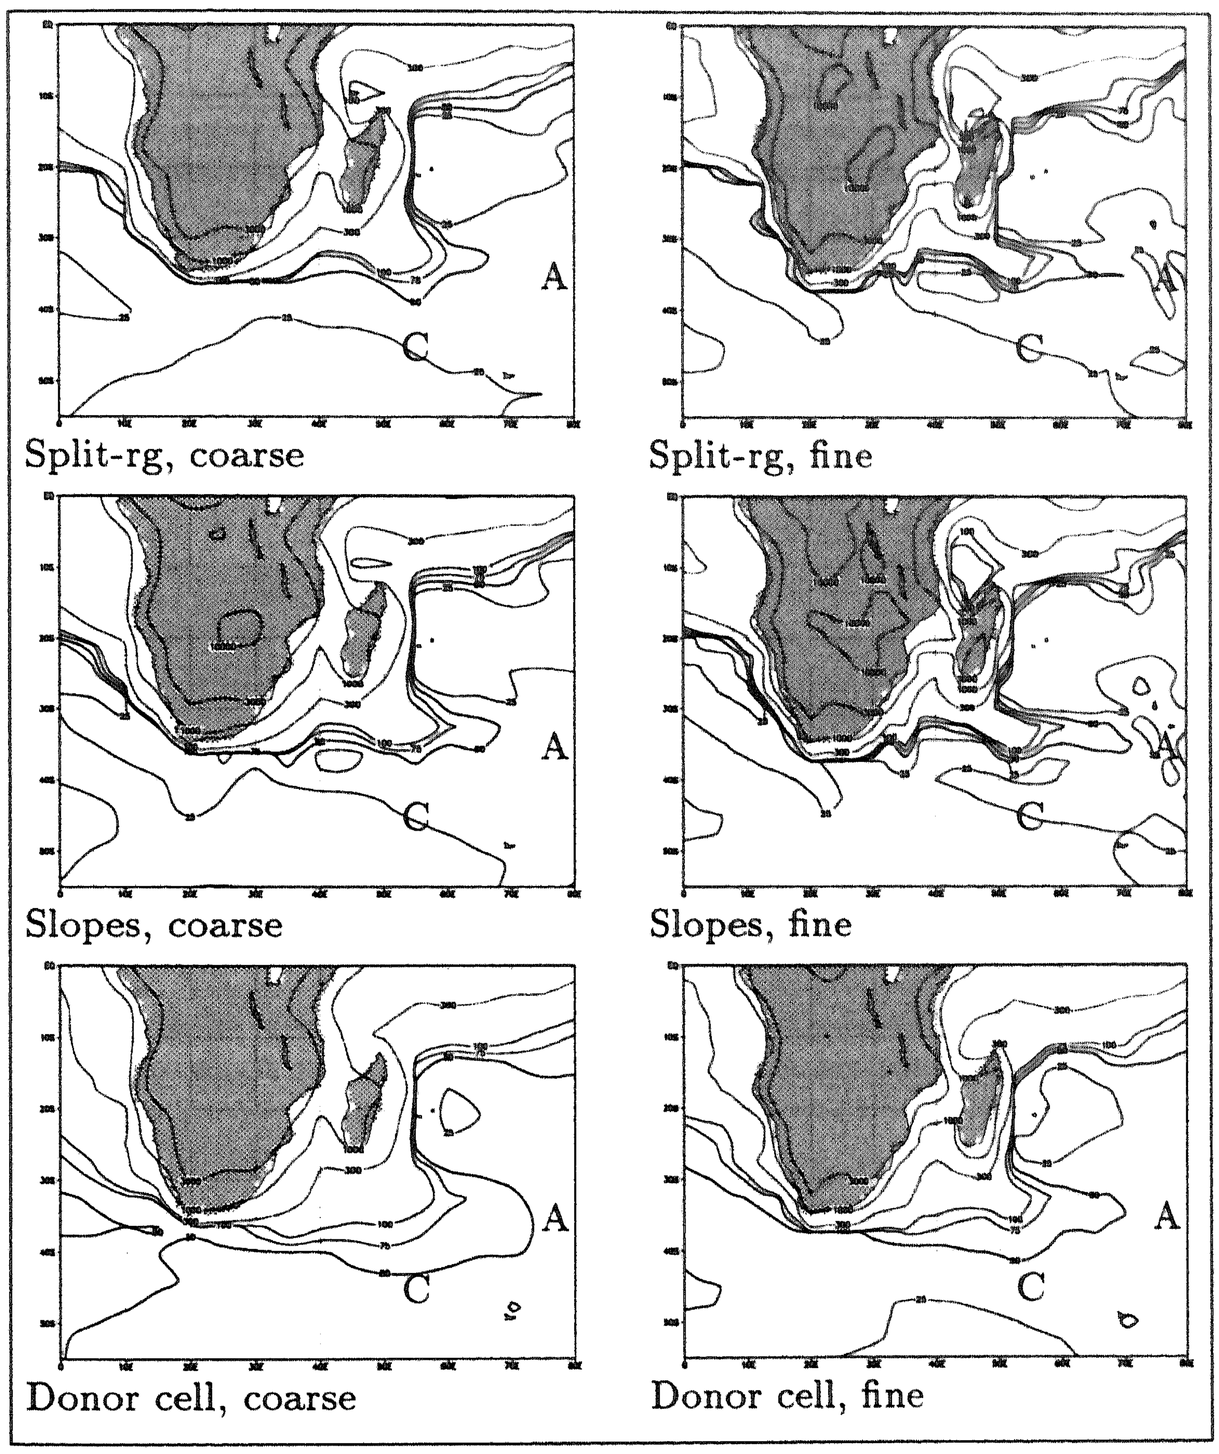

Figure 6.9: Contourlines of instantaneous ^{222}Rn concentration in mBq/SCM at 0000 UTC 21 January 1992 in a part of the Indian Ocean and southern Africa, for three advection schemes and two resolutions.

method	do_conv	calc_conv	advection	total
Donor Cell	11	4	10	26
Split-rg	10	4	18	33
Split-us	14	5	22	41
Mol-rg	11	4	41	57
Slopes	24	5	60	90
Second M.	52	5	147	205

Table 6.4: cpu times in minutes on a workstation for a simulation of one month on the $2.5° \times 2.5°$ grid for some parts of the transport model: the time integration of the convection, the calculation of the convection matrix, and the advection.

western-Africa, where the ^{222}Rn concentration is locally cone-shaped (not shown). From the solid-body rotation test we know that Slopes gives higher peak values than Split-rg. In the radon test this difference leads to about a 10% higher peak value in the 18-day average radon concentration field over the continents for Slopes. Since the differences between peak concentrations are confined to relatively small areas on the globe the continental-average near-surface concentrations differ approximately 4% between Split-rg and Slopes. The main purpose of this radon test for horizontal transport was to compare the performance of different numerical advection schemes under more realistic meteorological conditions. We have found that the numerical accuracy of Slopes, Second Moment and the four newly developed schemes is comparable for these conditions. Future studies on radon transport will have to take the following aspects into account: longer time integration, comparison between on-line and off-line transport model results, comparison with more measurement stations (including continental stations), and sensitivity to different sub-grid scale vertical transport parametrizations. A detailed investigation into some of these issues will be presented by Dentener et al. (Simulation of radon-222 using on-line and off-line global models, in preparation). The influence of different numerical advection schemes on the results will be very small, therefore the most efficient advection scheme can be used.

6.5.3 Performance results

In Table 6.4 results are given from a performance test, for a simulation of one month on the $2.5° \times 2.5°$ grid. We give results for the subroutines `do_conv`, `calc_conv` and `advection`. They respectively solve Equation (6.3), the sub-grid scale vertical transport, Equation (6.4), the calculation of the convection matrix, and the advection equation. The cpu time is in minutes on a workstation.

The time step for advection, Δt_{adv}, was always 30 minutes for the dimensional splitting schemes on the coarse grid, and depending on the actual wind velocities 15 or 30 minutes on the fine grid. The time step for Mol-rg and Donor Cell was 10, 15

or 30 minutes.

We see that on a reduced grid 15% cpu time is saved in `calc_conv`, and `do_conv` becomes expensive when using Slopes or Second Moment: it consumes more cpu time than Split-rg needs for advection. In a model with several chemical species the calculations in `do_conv` and `advection` must be done for all species involved, so it is important that these calculations are performed very efficiently.

From the radon test we know that the differences between Split-us, Split-rg, Mol-rg, Slopes and Second Moment are small, but from the results in Table 6.4 we see that Split-us and Split-rg are by far the most efficient solvers for pure advection, and all solvers on a reduced grid are 1.3 times faster in solving the sub-grid scale vertical transport parametrization than Split-us, three times faster than Slopes, and five times faster than Second Moment. We conclude that Split-us and Split-rg are the most efficient solvers for this application, but in models with extensive chemical schemes the reduced grid is recommended, because the chemistry is also solved 1.3 times faster and the efficiency on a vector processor of Split-us is lower.

6.6 Conclusion

We have presented four new advection schemes, especially designed for advection of trace gases in global air quality models. Important characteristics of these schemes are that they are mass conservative and that they do not generate undershoot and overshoot. Mass conservation is important because long-term calculations are performed with these models and the generation of undershoot may introduce negative values, possible resulting in instability when combined with stiff chemistry. The schemes are very efficient on (parallel) vector processors in terms of cpu time, memory requirements and accuracy. Also the pole singularity is solved.

We performed two tests: a solid-body rotation test, where we compared numerical results with an exact solution and a radon test where we compared model results with measurements. The new schemes are compared with two advection solvers that are currently widely used in global transport models: Slopes and Second Moment.

In the solid-body rotation test we showed that none of the new schemes generates undershoot and overshoot and that all of them are mass conservative. Split-us and Split-rg give accurate results and are very cheap in terms of cpu time and memory requirements. The schemes on a reduced grid, Split-rg and Mol-rg, and Split-us give good convergence speed, an important property when increasing the resolution. Unfortunately, Split-us is not very efficient on a vector processor.

Although in the radon test the general characteristics of the measurements were captured by the model results, we found large errors with respect to the height and timing of radon concentration peaks. Errors in the lower tropospheric flow patterns, due to a lack of meteorological measurements, as analyzed by the ECMWF for the month January 1992 in the Southern Indian Ocean region and errors related to the sub-grid scale vertical transport parametrizations must be considered as the main

sources of difference between the transport model and the measurements. However, the radon test also allows us to conclude that the differences between the advection schemes (here excluding the Donor Cell scheme) are small under realistic atmospheric conditions.

Slopes and Second Moment need a lot of cpu time in handling the moments during convection. They would have the same problem when chemistry is added to the model. The memory requirements for these schemes are respectively 4 and 10 times larger than for the other schemes.

At higher spatial resolutions the pole problem becomes more apparent. We have presented two options to overcome this difficulty: the reduced grid and the unconditionally stable Split-us scheme on a uniform grid. We generally recommend the use of Split-rg on a reduced grid because it has 25% less cells for the same spatial resolution at moderate latitudes, while good efficiency on a vector processor, such as the Cray C90, is obtained.

From both tests we conclude that Split-rg gives almost the same model results as Slopes and Second Moment for significantly less cpu time and memory requirements. Therefore we recommend Split-rg for solving the advection equation in global air quality models.

6.7 Appendix

6.7.1 Dimensional splitting in conservation form with mixing ratio fluxes

The new schemes presented and tested in this paper are 3D extensions of earlier derived schemes, such as Split-DeCo in [HS95], and the MOL schemes in [HKvLV95]. These schemes are based on fluxes as a function of the concentration. When the background density is not constant in space this approach does not give satisfactorily results. This is the case in the atmosphere because the density decreases with the altitude, and in the horizontal directions we have variations in the background density due to mountains and low and high pressure areas. However, good results are obtained if we use mixing ratio fluxes, which we will describe in this appendix. Even the accuracy problems with splitting schemes at the poles then disappear. We give formulas for first-order fluxes, but the extension to third-order fluxes is straightforward.

6.7.2 Splitting with concentrations

We consider the 3D advection equation in conservation form, which reads

$$c_t + (uc)_x + (vc)_y + (wc)_z = 0,$$

with unknown concentration, or vector of concentrations, c. A straightforward way to solve the advection equation numerically is by dimensional splitting. In its most

simple form this is

$$c^*_{i,j,k} = c^n_{i,j,k} + \left(F^n_{i-\frac{1}{2},j,k} - F^n_{i+\frac{1}{2},j,k}\right),$$

$$c^{**}_{i,j,k} = c^*_{i,j,k} + \left(G^*_{i,j-\frac{1}{2},k} - G^*_{i,j+\frac{1}{2},k}\right),$$

$$c^{n+1}_{i,j,k} = c^{**}_{i,j,k} + \left(H^{**}_{i,j,k-\frac{1}{2}} - H^{**}_{i,j,k+\frac{1}{2}}\right),$$

with concentration fluxes $F_{i+\frac{1}{2},j,k}$, $G_{i,j+\frac{1}{2},k}$ and $H_{i,j,k+\frac{1}{2}}$ computed by some 1D procedure, using interpolation of concentrations in x, y and z direction, respectively. For example, first-order upwind fluxes are given by

$$F_{i+\frac{1}{2}} = \begin{cases} \nu_{i+\frac{1}{2}} c_{i,j,k} \text{ if } u_{i+\frac{1}{2}} \geq 0, \\ -\nu_{i+\frac{1}{2}} c_{i+1,j,k} \text{ if } u_{i+\frac{1}{2}} < 0 \end{cases}$$

with $\nu_{i+\frac{1}{2}} = \left|u_{i+\frac{1}{2}}\right| \Delta t/\Delta x$ the local Courant number at the right cell boundary $\left(x_{i+\frac{1}{2}}, y_j, z_k\right)$.

Advection does not lead to new extremes. But even if the wind field is divergence free and the concentration is constant in space, a dimensional splitting scheme based on concentration fluxes can do just that. In the first step one tries to approximate the equation $c_t + (uc)_x = 0$. If, for example, we have inflow in both cell boundaries in the x direction of one grid cell, a new extreme can easily be created in that grid cell. The splitting scheme based on concentration fluxes does not guarantee that in the y and z step, this new extreme will vanish.

As observed by Bott [Bot92], first-order splitting may give qualitatively bad results with deformational flow fields. On the other hand, experiments in [HS95] suggested that for advection on a plane much better results are obtained if one uses a genuine second-order splitting method. Here the splitting should be second-order, for example Strang splitting [Str68], but also the 1D processes should be approximated with second-order accuracy. However, tests with advection on a sphere in [HS95] revealed that near the poles additional modifications were required, also with this second-order splitting.

The (2D) modification 'deformation correction' suggested in [HS95], consists of multiplying after each step the concentrations $c^{n+1}_{i,j}$ by a factor $\delta_{i,j}$ such that if we had started with $c^n_{i,j} \equiv 1$ then the $c^{n+1}_{i,j}$ would be the same as the first-order, non-splitted, Donor Cell algorithm applied to this uniform concentration field. Since this is a modification on the concentrations, not on the fluxes, the resulting scheme is not strictly mass conservative. Some tests in [HS95] on an analytical wind field of Williamson and Rasch [WR89] gave results which were 'almost' mass conserving. Still this point remained a matter of concern.

6.7.3 Splitting with mixing ratio fluxes

To overcome the creations of new extrema in the advection step, a modification on the fluxes was suggested by Russell and Lerner [RL81]. Our attention on this paper was drawn by the article of Easter [Eas92]. Also Allen et al. [ADRG91] proposed a similar approach.

Suppose the wind field is divergence free. Then a constant concentration field at time t_n should still be constant at t_{n+1}. Suppose now that we have given velocities that are divergence free in the discrete form

$$u_{i+\frac{1}{2}} - u_{i-\frac{1}{2}} + v_{j+\frac{1}{2}} - v_{j-\frac{1}{2}} + w_{k+\frac{1}{2}} - w_{k-\frac{1}{2}} = 0.$$

Then, if we have the densities $\rho^n_{i,j,k}$ at time t_n, and use these to compute the fluxes, the values at time t_{n+1} remain constant,

$$\rho^*_{i,j,k} = \rho^n_{i,j,k} + \frac{\tau}{\Delta x}\left(u_{i+\frac{1}{2}} - u_{i-\frac{1}{2}}\right),$$

$$\rho^{**}_{i,j,k} = \rho^*_{i,j,k} + \frac{\tau}{\Delta y}\left(v_{j+\frac{1}{2}} - v_{j-\frac{1}{2}}\right),$$

$$\rho^{n+1}_{i,j,k} = \rho^{**}_{i,j,k} + \frac{\tau}{\Delta z}\left(w_{k+\frac{1}{2}} - w_{k-\frac{1}{2}}\right) \equiv \rho^n_{i,j,k}.$$

This is in spite of the fact that the intermediate results $\rho^*_{i,j,k}$ and $\rho^{**}_{i,j,k}$ may give large variations. Note that the calculation of these intermediate results is similar to what is done in a splitting method.

The splitting modification of Russell and Lerner [RL81], consist of calculating the fluxes not from the concentrations $c^n_{i,j,k}$ and $c^*_{i,j,k}$, but from the mixing ratios $q_{i,j,k} = c_{i,j,k}/\rho_{i,j,k}$. The resulting scheme is

$$q^n_{i,j,k} = \frac{c^n_{i,j,k}}{\rho^n_{i,j,k}},\ c^*_{i,j,k} = c^n_{i,j,k} + \left(f^n_{i-\frac{1}{2},j,k} - f^n_{i+\frac{1}{2},j,k}\right),$$

$$q^*_{i,j,k} = \frac{c^*_{i,j,k}}{\rho^*_{i,j,k}},\ c^{**}_{i,j,k} = c^*_{i,j,k} + \left(g^*_{i,j-\frac{1}{2},k} - g^*_{i,j+\frac{1}{2},k}\right),$$

$$q^{**}_{i,j,k} = \frac{c^{**}_{i,j,k}}{\rho^{**}_{i,j,k}},\ c^{n+1}_{i,j,k} = c^{**}_{i,j,k} + \left(h^{**}_{i,j,k-\frac{1}{2}} - h^{**}_{i,j,k+\frac{1}{2}}\right),$$

with mixing ratio fluxes $f_{i+\frac{1}{2},j,k}$, $g_{i,j+\frac{1}{2},k}$ and $h_{i,j,k+\frac{1}{2}}$. These are computed by the same way as the concentration fluxes, only the values $q^n_{i,j,k}$, $q^*_{i,j,k}$ and $q^{**}_{i,j,k}$ are now used. For example, first-order upwind fluxes are given by

$$f_{i+\frac{1}{2}} = \begin{cases} \rho_{i+\frac{1}{2}}\,\nu_{i+\frac{1}{2}}\,q_{i,j,k} \text{ if } u_{i+\frac{1}{2}} \geq 0, \\ -\rho_{i+\frac{1}{2}}\,\nu_{i+\frac{1}{2}}\,q_{i+1,j,k} \text{ if } u_{i+\frac{1}{2}} < 0, \end{cases}$$

$$g^*_{j+\frac{1}{2}} = \begin{cases} \rho_{j+\frac{1}{2}}\,\nu_{j+\frac{1}{2}}\,q^*_{i,j,k} \text{ if } v_{j+\frac{1}{2}} \geq 0, \\ -\rho_{j+\frac{1}{2}}\nu_{j+\frac{1}{2}}\,q^*_{i,j+1,k} \text{ if } v_{j+\frac{1}{2}} < 0, \end{cases}$$

$$h^{**}_{k+\frac{1}{2}} = \begin{cases} \rho_{k+\frac{1}{2}}\,\nu_{k+\frac{1}{2}}\,q^{**}_{i,j,k} \text{ if } w_{k+\frac{1}{2}} \geq 0, \\ -\rho_{k+\frac{1}{2}}\nu_{k+\frac{1}{2}}\,q^{**}_{i,j,k+1} \text{ if } w_{k+\frac{1}{2}} < 0, \end{cases}$$

with $\nu_{j+\frac{1}{2}} = \left|v_{j+\frac{1}{2}}\right| \Delta t/\Delta y$ and $\nu_{k+\frac{1}{2}} = \left|w_{k+\frac{1}{2}}\right| \Delta t/\Delta z$ the local Courant numbers.

Suppose that the very first step of this algorithm gives $q^n_{i,j,k} = 1$ for all i, j, k. Then the interpolation for the fluxes is trivial, and we get $q^*_{i,j,k} = 1$ since the formulas for $c^{n+1}_{i,j,k}$ and $\rho^{n+1}_{i,j,k}$ will be the same. In a similar way it follows that $q^{n+1}_{i,j,k} = c^{n+1}_{i,j,k}/\rho^{n+1}_{i,j,k} \equiv 1$. This property is not shared by the original splitting. Due to the fact that the intermediate quantities $q^*_{i,j,k}$ may be far from equilibrium, the interpolation in the second step may give large errors in the original splitting.

The new splitting with third-order upwind fluxes with flux-limiting gives almost identical results for analytical wind field as Split-DeCo in [HS95], except for the fact that we now have strict mass conservation.

6.7.4 Convergence proof

It is not clear a priori whether splitting with the mixing ratio fluxes converges to the correct result. This will be demonstrated here for first-order upwind fluxes with velocities u, v positive, not necessarily divergence free, and with artificial densities $\rho^n_{i,j} \equiv 1$. We give this proof for 2D advection, the extension to 3D is straightforward. Then $q^n_{i,j} = c^n_{i,j}$ and

$$c^*_{i,j} = c^n_{i,j} + \frac{\tau}{\Delta x}\left(u_{i-\frac{1}{2},j}c^n_{i-1,j} - u_{i+\frac{1}{2},j}c^n_{i,j}\right),$$

$$q^*_{i,j} = c^*_{i,j}/\rho^*_{i,j} \text{ with } \rho^*_{i,j} = 1 + \frac{\tau}{\Delta x}\left(u_{i-\frac{1}{2}} - u_{i+\frac{1}{2}}\right),$$

$$c^{n+1}_{i,j} = c^*_{i,j} + \frac{\tau}{\Delta y}\left(v_{j-\frac{1}{2}}q^*_{i,j-1} - v_{j+\frac{1}{2}}q^*_{i,j}\right).$$

This last formula can also be written as

$$c^{n+1}_{i,j} = c^*_{i,j} + \frac{\tau}{\Delta y}\left(v^*_{j-\frac{1}{2}}c^*_{i,j-1} - v^*_{j+\frac{1}{2}}c^*_{i,j}\right),$$

with $v^*_{j+\frac{1}{2}} = v_{j+\frac{1}{2}}/\rho^*_{i,j}$. Note that

$$\rho^*_{i,j} = 1 + \frac{\tau}{\Delta x}\left(u_{i-\frac{1}{2},j} - u_{i+\frac{1}{2},j}\right) \approx 1 + \tau u_x(x_i, y_j).$$

So, the above procedure is equivalent with first-order upwind splitting in terms of the concentrations, but with a modified velocity

$$u^* = u,\ v^* = \frac{v}{1 + \tau(u_x + \mathcal{O}(\Delta x))} = (1 + \mathcal{O}(\tau))\ v.$$

First-order upwind splitting with these velocities will approximate the exact solution of $c_t^* = (u^*c^*)_x + (v^*c^*)_y$ with a global error $\mathcal{O}(h)$, where $h = \max(\Delta x, \Delta y)$. Further the difference of $c(x, y, t)$ and $c^*(x, y, t)$ is of $\mathcal{O}(\tau)$. Thus the above scheme approximates $c(x, y, t)$ with an error $\mathcal{O}(\tau) + \mathcal{O}(h)$. Note that it has not been assumed that the wind field is divergence free.

Chapter 7

A second-order Rosenbrock method applied to photochemical dispersion problems

Abstract

A second-order, L-stable Rosenbrock method from the field of stiff ordinary differential equations is studied for application to atmospheric dispersion problems describing photochemistry, advective and turbulent diffusive transport. Partial differential equation problems of this type occur in the field of air pollution modeling. The focal point of this chapter is to examine the Rosenbrock method for reliable and efficient use as an atmospheric chemical kinetics box-model solver within Strang-type operator splitting. In addition two W-method versions of the Rosenbrock method are discussed. These versions use an inexact Jacobian matrix and are meant to provide alternatives for Strang-splitting. Another alternative for Strang-splitting is a technique based on so-called source-splitting. This technique is briefly discussed.

Note: This chapter is based on report MAS-R9717, and has been submitted for publication in *SIAM Journal on Scientific Computing*.

7.1 Introduction

Photochemical dispersion models are used to enhance the understanding of the chemical composition of the atmosphere, in particular with regard to the relation between

anthropogenic emissions and the resulting distributions of primary and secondary polluting species. Modern models are based on mass balances in the form of systems of time-dependent, three-space dimensional, partial differential equations (PDEs) describing advective transport, turbulent diffusive transport, cumulus cloud convection, chemical reactions, emissions and depositions. Models are discretized on Eulerian grids over areas of different sizes, from urban to regional to fully global. Comprehensive dispersion models are computationally extremely expensive. After spatial discretization they lead to huge systems of ordinary differential equations (ODEs) containing stiff nonlinear terms from the photochemical reactions [Zla95]. For the time integration of these ODEs one commonly uses Strang-type operator splitting, distinguishing mainly between transport and chemistry. The chemistry computation then amounts to repeatedly solving box-models $\dot{c} = f(c)$ at any of the spatial grid points. For this purpose dedicated explicit methods are used [HBCH93, JSPC97, Ver94, VS95] or solvers from the stiff ODE field adjusted for exploiting sparsity in the Jacobian matrix [SPCD96, JT94, SVvL$^+$97, SVB$^+$97, VBvLS96]. As a rule, black-box solvers are considered too costly.

The accuracy demand for atmospheric dispersion problems is modest so that a low order method is a natural choice. Rosenbrock methods have already proven very effective for low to modest accuracies for a wide variety of stiff problems [HW96]. When sparsity is exploited, their specific advantage for integrating atmospheric box-models has recently been demonstrated in [SVB$^+$97], where the stiffly accurate solver RODAS3 came out as most competitive. This solver is a variable step size, third-order, 4-stage counterpart of the well-known fourth-order, 6-stage solver RODAS from [HW96]. However, like most solvers, RODAS3 is quite sensitive to initial transients and therefore can require a rather small step size in the initial phase to start up the integration. For single ODE systems and long integration intervals, a small initial step renders no problem as it hardly adds to the total work load. In the context of operator splitting the situation is different. The integration intervals are relatively short and usually small transients are encountered within any split step. These transients are insignificant and for efficiency reasons it is desirable to start up with an a priori described step size far greater than the smallest time constant. In addition, when running operator splitting schemes on vector/parallel or massively parallel computers, a priori described step sizes everywhere seem more practical than truly variable ones. Variable step size integration governed by local error control creates load imbalancing in parallel implementations [Elb97] and interferes with vectorization over the horizontal grid dimension [JT94, SVdZ$^+$97a].

These observations have led us to search for a more stable Rosenbrock formula which is capable of integrating with much larger a priori described step sizes. The result is a particular, L-stable, second-order 2-stage method. In the remainder this method will be called ROS2 (Rosenbrock, second-order). The main difference between ROS2 and RODAS3 lies in the stability function and the internal stability functions. For ROS2 both these functions are positive along the negative real line, like the

exponential. This does not hold for RODAS3. In this chapter we provide strong numerical evidence that this simple positivity property significantly enhances the nonlinear stability for atmospheric chemical kinetic problems, thus rendering ROS2 highly efficient for use within Strang-type operator splitting. Here lies the focal point of this chapter. Another difference with RODAS3 is that for second-order consistency, ROS2 can be used with an inexact Jacobian matrix. This property can be exploited in different ways and we pay some attention to it in connection with alternatives for standard Strang-splitting. We note in passing that Rosenbrock methods using an inexact Jacobian are also called W-methods [HW96].

The contents of this chapter is as follows. In Section 7.2 we outline the main intended application for ROS2 through a prototype model for 3D spherical photochemical dispersion. This prototype model will be used as a test problem along with three different sets of atmospheric photochemical reactions from the actual practice. In Section 7.3 we present ROS2, discuss the basic properties which render this Rosenbrock method suitable and highly competitive as an atmospheric box-model integrator and provide numerical results to illustrate this. In Section 7.4 we apply ROS2 within the context of Strang-splitting. Within splitting chemistry is decoupled from transport, yielding a potential difficulty for the chemistry integration in that stiff transients can arise in any split step. These transients are artefacts of the splitting and require a robust and stable method. The main objective of this section is show that ROS2 is well capable for this task. Two different second-order Strang-splittings are used. The first treats advection and diffusion explicit and coupled, the second treats advection explicit and vertical diffusion implicit, but still coupled. At this stage we introduce a W-version of ROS2 as an advection-diffusion solver. We use this solver within Strang-splitting, but it is of obvious interest of its own. Section 7.5 is devoted to a complete alternative for Strang-splitting. Here we present a second application of ROS2 as a W-method, but now for the full advection-diffusion-chemistry problem. The approximate factorization approach is used here for the inexact Jacobian definition. This approximate factorization ROS2 scheme has been introduced to provide a comparison with Strang-splitting for cases where there is large vertical turbulent diffusion. Standard Strang-splitting then yields larger splitting errors and comparison with alternatives is of numerical interest. The result of the present comparison is that the approximate factorization ROS2 scheme certainly is competitive, but since Strang-splitting is somewhat simpler for use in the actual practice, greater benefits must be shown to replace it. In Section 7.6 we briefly discuss another alternative for Strang-splitting, which is called source-splitting. However, on theoretical grounds it is argued that Strang-splitting is more stable. The main conclusion of our investigations is formulated in the final Section 7.7.

7.2 Photochemical dispersion models

By way of illustration we first outline the spherical, global prototype model from [SVdZ+97a] which later will serve as a test model. Let $c = c(t, \lambda, \phi, r)$ denote a vector of N_c species concentrations with t representing the time, $\lambda \in [0, 2\pi]$ the longitude, $\phi \in [-\frac{\pi}{2}, +\frac{\pi}{2}]$ the latitude and $0 \leq r \leq r_H$ the height above the surface of the earth. Let a be the radius of the earth, ρ the density of the air, (u, v) a horizontal wind field and K a vertical subgrid-scale turbulent diffusion coefficient. Let $f(c)$ be an N_c-dimensional vector function representing atmospheric (photo)chemical reactions, emission sources and sink depositions. The mathematical formulation of the prototype model then reads

$$\frac{\partial c}{\partial t} + \frac{1}{a \cos\phi}\left[\frac{\partial (u\, c)}{\partial \lambda} + \frac{\partial (v\, c \cos\phi)}{\partial \phi}\right] = \frac{\partial}{\partial r}\left(\rho K \frac{\partial}{\partial r}\left(\frac{c}{\rho}\right)\right) + f(c). \qquad (7.1)$$

The reaction term $f(c)$ couples the N_c equations in the model. Without reactions the equations are not coupled since u, v, ρ and K are given expressions. Processes are time and space dependent. In real models the various meteorological parameters are kept constant in time and updated every few hours, say. For numerical purposes we thus may consider the processes constant in time. At the surface and the top the no-flux boundary conditions

$$\rho K \frac{\partial}{\partial r}\left(\frac{c}{\rho}\right) = 0 \qquad (7.2)$$

are imposed. The model is completed by prescribing the initial values for c at the initial time $t = t_0$. The unit for the concentrations is number of molecules per cm^3 (mlc/cm^3). The unit of time is seconds.

7.2.1 Transport

In the prototype model the wind field (u, v), the density ρ and the coefficient K are given analytic expressions. The wind field is a solid body rotation with a maximum speed of about 125 km/hour. The diffusion coefficient K is dependent on the height and takes on a maximum of 30 m^2/s in the lower troposphere. We used a parameterization given in [War88, page 24, Figure 1-10]. The top of the model lies at 34.7 km. Real models also simulate transport by vertical advection, subgrid-scale turbulent horizontal diffusion and subgrid-scale cumulus cloud convection. For a numerical study horizontal diffusion is not really essential. Horizontal diffusion can always be added and numerically integrated explicitly, in a similar manner as horizontal advection. In a same manner, vertical advection can often be treated explicitly. If not, it can be combined somehow with the implicit vertical diffusion computation. Whether the absence of cloud convection is essential is as yet not clear. In reality there is also orography. This makes models technically much more complicated as this amounts to a coordinate transformation of the ideal sphere model (7.1). We believe, however, that with regard to time stepping the absence of orography is not essential either.

7.2.2 Chemistry

Atmospheric photochemistry induces severe stiffness. Reaction times may range from milliseconds or shorter (e.g. OH radical) to years (e.g. CH_4). The photochemical nature complicates the numerical solution, since part of the reaction coefficients depends on the solar zenith angle which depends on the time of the day and the location on earth. This dependence gives rise to constantly moving areas of rapid solution change coupled to sunset and sunrise. There also exists a dependence on the temperature and the pressure. This dependence is chosen in close accordance with the US Standard Atmosphere (1976) [War88]. In applications the number of species varies, between 20 and 100 say. In numerical illustrations presented further on we use three different sets of chemical reactions, all borrowed from the actual practice:

- Chemistry model RIVM: The first set consists of 45 reactions between $N_c = 17$ species and is used in actual long term global studies where it is referred to as methane chemistry. See the appendix of the preprint of [VBvLS96].

- Chemistry model CBM-IV: The second set is based on the Carbon Bond Mechanism IV consisting of $N_c = 32$ chemical species involved in 70 thermal and 11 photolytic reactions. We also used this model in [SVB+97], but with a different solar zenith angle. To stress the numerical method, we have prescribed high emission values (the urban scenario from [SVB+97]).

- Chemistry model WET: The third set contains $N_c = 65$ species involved in 77 thermal and 11 photolytic gas-phase chemical reactions, 39 liquid-phase chemical reactions and 39 gas-liquid mass transfer reactions. The gas-phase mechanism is based on the Carbon Bond Mechanism IV, while the liquid-phase mechanism is based on a chemical scheme from [Mat95]. We also used this model in [SVB+97], again with a different solar zenith angle, and emphasize that it is the most difficult one from the three mentioned here due to the heterogeneous reactions.

7.2.3 Spatial discretization

Model (7.1) is discretized on a 3D Eulerian grid spanning the entire globe. The longitude-latitude grid is uniform, except near the poles where the grid is reduced (coarsened) in the longitude direction to alleviate the CFL restriction for explicit advection schemes caused by the pole singularity. The spatial advection scheme is based on a mass-conservative, cell-centered, flux-limited, third-order upwind discretization. Flux limiting is used for positivity. See [HKvLV95] and [SVdZ+97a] for details.

The spatial vertical diffusion scheme is based on cell-centered 3-point finite-differences. The vertical grid is nonuniform. The prototype model has 15 layers. The distribution of the cell vertical centers is a function of the pressure which is taken uniform over the globe. The lowest cell boundary lies at sea level (1000 hPa) and the

highest at 38.2 km (0 hPa). The complete distribution of the cell centers reads 0.3, 1.0, 2.2, 4.3, 6.5, 8.4, 10.0, 11.3, 13.0, 15.2, 17.6, 19.8, 22.5, 27.6, 34.7 km.

In [SVdZ+97a] three different reduced longitude-latitude grids were used. Without grid reduction their dimensions are 64×32, 128×64 and 256×128, respectively. With grid reduction this leads to 1656, 6264 and 26104 horizontal grid points, so that combined with the vertical grid totals of 24840, 93960 and 391560 grid points were used. Recall that for any grid point and any of the N_c species, a nonlinear ODE must be integrated in time, revealing the enormous computational scale of atmospheric dispersion modeling and the necessity of developing highly efficient, tailored algorithms. See also the recent monograph of Zlatev [Zla95].

7.3 Solving box-models

In atmospheric dispersion modeling one frequently applies operator splitting and employs stiff ODE solvers to integrate resulting box-models. In this section we therefore first consider the box-model

$$\dot{c} = f(c) \tag{7.3}$$

contained in (7.1). We will outline our specific choice of Rosenbrock method, discuss the basic properties which renders this method suitable and highly competitive as an atmospheric box-model integrator and provide numerical results to illustrate this.

7.3.1 The ROS2 integration formula

Our starting point is a family of nonautonomous 2-stage Rosenbrock methods discussed in [DV84, page 233]:

$$\begin{aligned} c_{n+1} &= c_n + \tau b_1 k_1 + \tau b_2 k_2, \\ k_1 &= f(t_n, c_n) + \tau\gamma A\, k_1, \\ k_2 &= f(t_n + \tau\alpha_{21}, c_n + \tau\alpha_{21} k_1) + \tau\gamma_{21} A\, k_1 + \tau\gamma A\, k_2, \end{aligned} \tag{7.4}$$

where $c_n \approx c(t)$ at $t = t_n$, $\tau = t_{n+1} - t_n$ is the step size and A is the Jacobian matrix $f'(t_n, c_n)$ or an approximation thereof. The method is second-order consistent for any A iff

$$b_1 = 1 - b_2, \quad \gamma_{21} = -\gamma/b_2, \quad \alpha_{21} = 1/(2b_2),$$

with γ and $b_2 \neq 0$ still free. We will use the autonomous form and in this section we assume that $A = f'(c_n)$ which ensures conservation of mass. The parameter γ appears in the stability function

$$R(z) = \frac{1 + (1 - 2\gamma)z + (\frac{1}{2} - 2\gamma + \gamma^2)z^2}{(1 - \gamma z)^2}. \tag{7.5}$$

This function is A-stable iff $\gamma \geq 1/4$. Since atmospheric chemistry models contain radicals with a very short life span we want L-stability, that is, $R(\infty) = 0$. This is achieved by $\gamma = 1 \pm 1/\sqrt{2}$. Further we select $b_2 = \frac{1}{2}$. Avoiding the matrix-vector multiplication in the second stage computation, and at the same time redefining k_2 by $k_2 - 2k_1$, the resulting autonomous ROS2 scheme then is rewritten in the form

$$\begin{aligned} c_{n+1} &= c_n + \tfrac{3}{2}\tau k_1 + \tfrac{1}{2}\tau k_2, \\ (I - \gamma\tau A)\, k_1 &= f(c_n), \\ (I - \gamma\tau A)\, k_2 &= f(c_n + \tau k_1) - 2k_1. \end{aligned} \tag{7.6}$$

Observe that the intermediate approximation $c_n + \tau k_1$ is 1st-order consistent at $t = t_{n+1}$ and hence can be used to provide a cheap local error estimation for step size control. In the present investigation we will not exploit this possibility since we focus on using a priori described step sizes as mentioned in the introduction.

7.3.2 Stability and positivity

For γ we select the larger value $\gamma_+ = 1+1/\sqrt{2}$ in spite of the fact that this gives a larger error coefficient in the leading local truncation error. Numerical tests have revealed that this yields a notably better nonlinear stability behaviour for large step sizes. A balanced explanation fails, but we conjecture that the following linear property plays a role. If $\gamma = \gamma_+$, then $R(z)$ is positive for all real negative z, whereas this is not true for the smaller value $\gamma_- = 1 - 1/\sqrt{2}$. Positivity of R presumably has some advantage for nonlinear chemical kinetic systems

$$\dot{c_k} = f_k(c) \equiv P_k(c) - L_k(c)c_k,$$

where $P_k(c)$ contains all production terms for the $k-$th species and $L_k(c)c_k$ represents the losses for this species. Suppose that for a certain species, P_k and L_k are truly constant. Then,

$$c_{k,n+1} = R(z)\, c_{k,n} + \frac{R(z) - 1}{z}\tau P_k, \quad z = -\tau L_k \leq 0.$$

If $R(z) < 0$, then $c_{k,n+1}$ might become negative. On the other hand, if $0 \leq R(z) \leq 1$, then $c_{k,n+1} \geq 0$ is guaranteed. This obviously proves nothing for truly nonlinear systems. However, in the atmosphere the so-called radicals react very fast and are always near to their steady state value $P_k(c)/L_k(c)$. If the dependence of P_k, L_k on c is very weak, the above linear reasoning can come close to what happens in the Rosenbrock computation. Even a very small negative solution value can cause sign problems, because radicals occur in a multiplication with very large positive reaction constants. Another advantage of γ_+ is that the internal stability function

$$R_1(z) = \frac{1 + (1 - \gamma)z}{1 - \gamma z} \tag{7.7}$$

associated to the first-stage approximation $c_n + \tau k_1$ also satisfies $0 < R_1(z) < 1$ for all real negative z, something which does not hold for γ_-. This property of internal stability has already been shown to be of practical interest for nonlinear stiff ODE problems in [Ver77].

Of further interest is that γ_+ is also to be preferred when solving the nonlinear scalar model problem

$$\dot{c} = \lambda c^2, \quad \lambda < 0, \tag{7.8}$$

whose solution

$$c(t) = \frac{c(0)}{1 - \lambda c(0) t}$$

remains positive if $c(0)$ is positive. Denote $z = \tau\lambda$. The approximations $c_n + \tau k_1$ and c_{n+1} then read

$$c_n + \tau k_1 = \frac{c_n + (1 - 2\gamma) z c_n^2}{1 - 2\gamma z c_n}$$

and

$$c_{n+1} = \frac{c_n + (1 - 6\gamma) z c_n^2 + (1 - 6\gamma + 12\gamma^2) z^2 c_n^3 + (\frac{1}{2} - 2\gamma + 8\gamma^2 - 8\gamma^3) z^3 c_n^4}{(1 - 2\gamma z c_n)^3}.$$

It easily follows that both are unconditionally positive for $\gamma = \gamma_+$, whereas both can become negative if $\gamma = \gamma_-$. For the intermediate approximation $c_n + \tau k_1$, the term $(1 - 2\gamma) z c_n^2$ in the numerator fails to be unconditionally positive and for c_{n+1} this is the case for the term $(\frac{1}{2} - 2\gamma + 8\gamma^2 - 8\gamma^3) z^3 c_n^4$.

7.3.3 Clipping

For real photochemical systems positivity cannot be guaranteed. Because it is essential for stability, in the application of ROS2 positivity is enforced at both the stages by clipping. This means that when a component of $c_n + \tau k_1$ or c_{n+1} is negative, it is set equal to zero. Clipping interferes with the property of mass conservation. However, in our experiments we have not observed a notable loss in accuracy, presumably because in an actual integration ROS2 is clipping only occasionally if $\gamma = \gamma_+$. In Section 7.3.5 we will provide numerical evidence for this observation. Lest we miss the obvious, enforcing positivity by clipping does not guarantee stability.

7.3.4 Workload and sparsity

Each time step with ROS2 requires an evaluation of the Jacobian $A = f'(c_n)$, two linear system solutions accompanied with two derivative evaluations. The Jacobian update and the solution of the linear systems, requiring one matrix factorization (LU-decomposition) and two backsolves (forward-backward substitutions), account for most of the CPU time. Fortunately, for large atmospheric chemistry models the

Chemical reaction set	WET	CBM-IV	RIVM
Number of entries in F'	65^2	32^2	17^2
Number of nonzeroes in F'	506	276	100
Number of nonzeroes after sparse LU	629	300	107

Table 7.1: Sparsity data for the three sets of chemical reactions.

number of zeroes in the Jacobian is substantial. For very large models it readily amounts to $\approx 90\%$. This high level of sparsity can be exploited to significantly reduce the costs of these linear algebra calculations. For this purpose we use the symbolic preprocessor KPP [DIS95, SPCD96] in the same way as in [SVB+97]. KPP takes as input a set of chemical reactions and delivers the production and loss terms defining the ODE system $\dot{c} = f(c)$. Most important is that it also prepares a sparse matrix factorization with only a minimal fill-in and that it delivers a routine for the backsolve without indirect addressing. Altogether this means that the numerical algebra can be handled very efficiently, leading to a substantial reduction of the workload for large chemical kinetic models. Table 7.1 shows sparsity data for the three models WET, CBM-IV and RIVM.

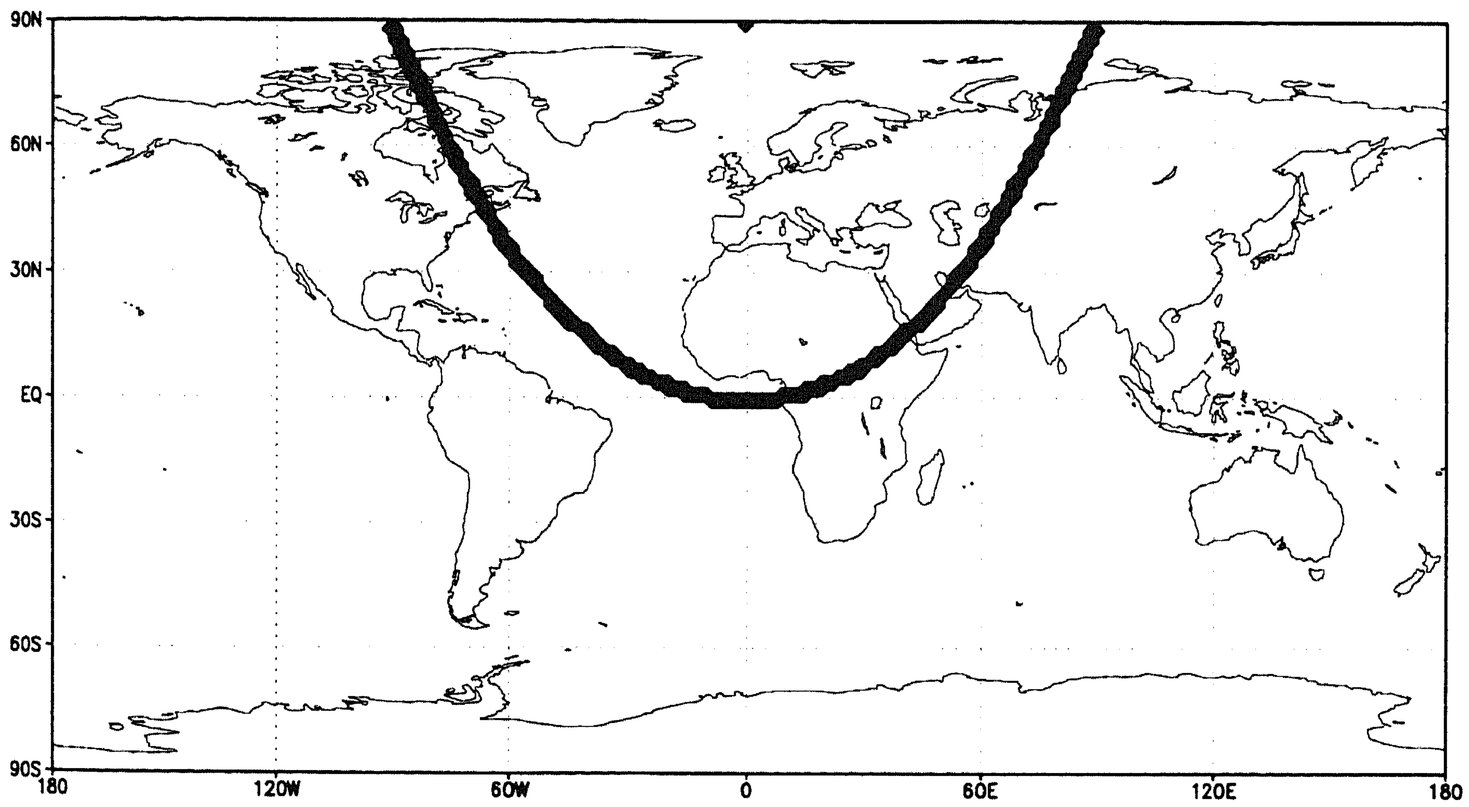

Figure 7.1: The trajectory used in the box-model tests and 1D tests.

7.3.5 Numerical illustrations

Numerical results will be shown for three box-models based on the reaction sets RIVM, CBM-IV and WET. Starting from an arbitrary initial state far from chemical equilibrium, in all tests we simulate that we follow an air parcel from its release point east of Africa for 14 days along the trajectory shown in Figure 7.1, picking up emissions along the way. After 14 days the air parcel returns at its release point. Because to a great extent the initial conditions are chosen arbitrary, the first day is used as start up time (integration over the first day is carried out nearly exact using a very small step size). It should be noted that the trajectory passes the North Pole at day 7, in the neighbourhood of which the photochemical reactions are weaker than elsewhere. This effect leads to a disturbance in the diurnal behaviour which can be observed in the species solutions. Hourly frozen reaction coefficients were used with an update half way each hour interval. This renders the ODE systems autonomous and implies that the pressure, temperature and solar angle are taken piecewise constant rather than time-continuous. In the actual practice one normally operates this way, one reason being that many of the coefficients are expensive to compute.

We first present results of a stability test, comparing ROS2 for $\gamma = \gamma_+$ and $\gamma = \gamma_-$. We have also included the related Rosenbrock method RODAS3 proposed in [SVB+97] in this test. RODAS3 is based on a stiffly accurate, embedded pair of order 3(2). It has 4 stages and uses 3 function evaluations. In our stability test we have used only the third-order formula which is given by

$$\begin{aligned} c_{n+1} &= c_n + \tfrac{5}{6}\tau k_1 - \tfrac{1}{6}\tau k_2 - \tfrac{1}{6}\tau k_3 + \tfrac{1}{2}\tau k_4, \\ (I - \tfrac{1}{2}\tau A)\, k_1 &= f(c_n), \\ (I - \tfrac{1}{2}\tau A)\, k_2 &= f(c_n) + Ak_1, \\ (I - \tfrac{1}{2}\tau A)\, k_3 &= f(c_n + \tau k_1) - \tfrac{1}{4}Ak_1 - \tfrac{1}{4}Ak_2, \\ (I - \tfrac{1}{2}\tau A)\, k_4 &= f(c_n + \tfrac{3}{4}\tau k_1 - \tfrac{1}{4}\tau k_2 + \tfrac{1}{2}\tau k_3) + \tfrac{1}{12}Ak_1 + \tfrac{1}{12}Ak_2 - \tfrac{2}{3}Ak_3. \end{aligned} \tag{7.9}$$

As already mentioned in the introduction, in actual integrations we have experienced that this third-order, 4-stage method is less stable than ROS2 when using fixed, large step sizes. As for ROS2 using $\gamma = \gamma_-$, we conjecture that lack of positivity of its stability function

$$R(z) = \frac{1 - z + \frac{1}{6}z^3}{(1 - \frac{1}{2}z)^4}$$

plays a role here. Further, from its three internal stability functions

$$R_1(z) = 1, \quad R_2(z) = \frac{1 + \frac{1}{2}z}{1 - \frac{1}{2}z}, \quad R_3(z) = \frac{1 - \frac{1}{2}z - \frac{1}{4}z^2}{(1 - \frac{1}{2}z)^3},$$

both R_2 and R_3 fail to be positive.

Table 7.2 shows maximal step sizes for which integrations were found stable. The integrations cover 13 days and start at day 2 from the chemical equilibrium as outlined

above. The table gives clear evidence of the superior stability of ROS2 for the larger γ. Starting from the arbitrary initial states of day 1 far away from the chemical equilibrium shows even greater differences. See the italic numbers in Table 7.2. The table also reveals that the result of clipping may lead to a smaller step size rather than to a larger one, as happens for ROS2(γ_+) applied to RIVM. However, in all other cases clipping indeed does allow larger step sizes. Of further interest is that with respect to stability, RODAS3 performs notably better than ROS2(γ_-), but less than ROS2(γ_+). The difference between the latter two is much less though. We believe that RODAS3 owes this to its higher order of consistency which results in more accuracy. More accuracy will eventually lead to positive solutions and hence to a more stable process.

For this specific test an integration has been called stable if during the whole integration period a certain relative error remains below 10. The precise error definition is not so important, nor is the threshold 10. We used the error expression

$$\frac{1}{N_c}\sum_{k=1}^{N_c} ER_k,$$

where

$$ER_k = \sqrt{\frac{1}{|\mathcal{J}_k|}\cdot\sum_{n\in\mathcal{J}_k}\left|\frac{c_k(t_n)-\hat{c}_k(t_n)}{\hat{c}_k(t_n)}\right|^2},$$

and

$$\mathcal{J}_k = \{0\le n\le N : \hat{c}_k(t_n)\ge a_k\},\quad a_k = \frac{10^{-4}}{N}\sum_{i=1}^{N}\hat{c}_k(t_i).$$

The solution values $\hat{c}_k(t_n)$ represent accurate reference solutions at every full hour value t_n. The set $\mathcal{J}_k$ has been introduced to remove very small solution values in the relative measurement.

We proceed with presenting results obtained for the actual box-model integrations by ROS2 and TWOSTEP. The latter is a two-step BDF solver using nonlinear Gauss-Seidel iteration instead of modified Newton for the nonlinear BDF relations [Ver94]. The use of Gauss-Seidel iteration renders this solver effectively explicit and hence cheap. TWOSTEP is capable of solving gas-phase chemistry more efficiently than other dedicated explicit solvers, e.g. QSSA [JSPC97, SVvL+97]. But like all other explicit solvers it cannot handle heterogeneous reactions as in WET. Therefore, results for TWOSTEP only concern the models RIVM and CBM-IV (see also [SVB+97]).

For the two species O_3 and HNO_3, Figures 7.2 and 7.3 show concentrations in mlc/cm^3 versus time in hours for day 2 up to day 14. The figures contain a highly accurate reference solution and the two computed solutions, using a fixed step size of 10 min and 20 min, respectively. Such fixed step sizes are very large for atmospheric chemistry integrations. Observe that in many applications advection step sizes are also in this range.

Method	WET	CBM-IV	RIVM
ROS2 (γ_+)	3600 - 3600	3600 - 1800	1200 - 1800
	3600 - 1800	*3600 - 150*	*1200 - 1800*
ROS2 (γ_-)	150 - 90	212 - 156	400 - 200
	< 5 - < 5	*133 - 36*	*400 - 133*
RODAS3	1800 - 400	1800 - 514	1200 - 900
	1800 - <5	*1800 - 133*	*1200 - 900*

Table 7.2: Box-model stability test. The numbers are step sizes in seconds with - without clipping. Italic numbers belong to the start point outside chemical equilibrium.

The Rosenbrock method can be seen to perform very satisfactorily. With the smallest step size it computes O_3 and HNO_3 for WET up to plotting accuracy. For $\tau = 20$ min a mild error growth occurs for this model. Applied to CBM-IV and RIVM, the Rosenbrock method delivers excellent results for both step sizes. TWOSTEP also solves CBM-IV accurately with the smallest step size, but generates significant errors for $\tau = 20$ min. When applied to RIVM this explicit code clearly requires a smaller step size than the Rosenbrock method.

Noteworthy is that the accuracy for TWOSTEP can be improved by spending more Gauss-Seidel iterations. Only a fixed number of 2 iterations has been used here, similar as in [SVB+97]. This makes it very cheap in CPU. The CPU time needed by ROS2 for CBM-IV and RIVM is only about a factor 2 higher, revealing the efficiency of this linearly implicit solver. In part we owe this to the use of KPP and the sparsity. However, the main advantage of ROS2 over TWOSTEP and related dedicated explicit solvers is that it can deal with different kinds of reactions, including heterogeneous ones as in WET.

7.4 ROS2 within Strang-type operator splitting

The Rosenbrock solver ROS2 is primarily meant for efficient use in standard, Strang-type operator splitting codes. In this section we will illustrate that the solver is capable for this task by showing results for the prototype model from Section 7.2. Our findings obviously also apply to different models. Adopting the method of lines approach, let

$$\dot{w} = F(w) \equiv F_T(w) + F_R(w), \quad F_T(w) \equiv F_A(w) + F_D(w) \tag{7.10}$$

denote the ODE system that originates from spatial discretization of the photochemical transport model. Hence $w(t)$ stands for a grid function and the vector function F_T is supposed to contain the semi-discrete transport contributions from advection and

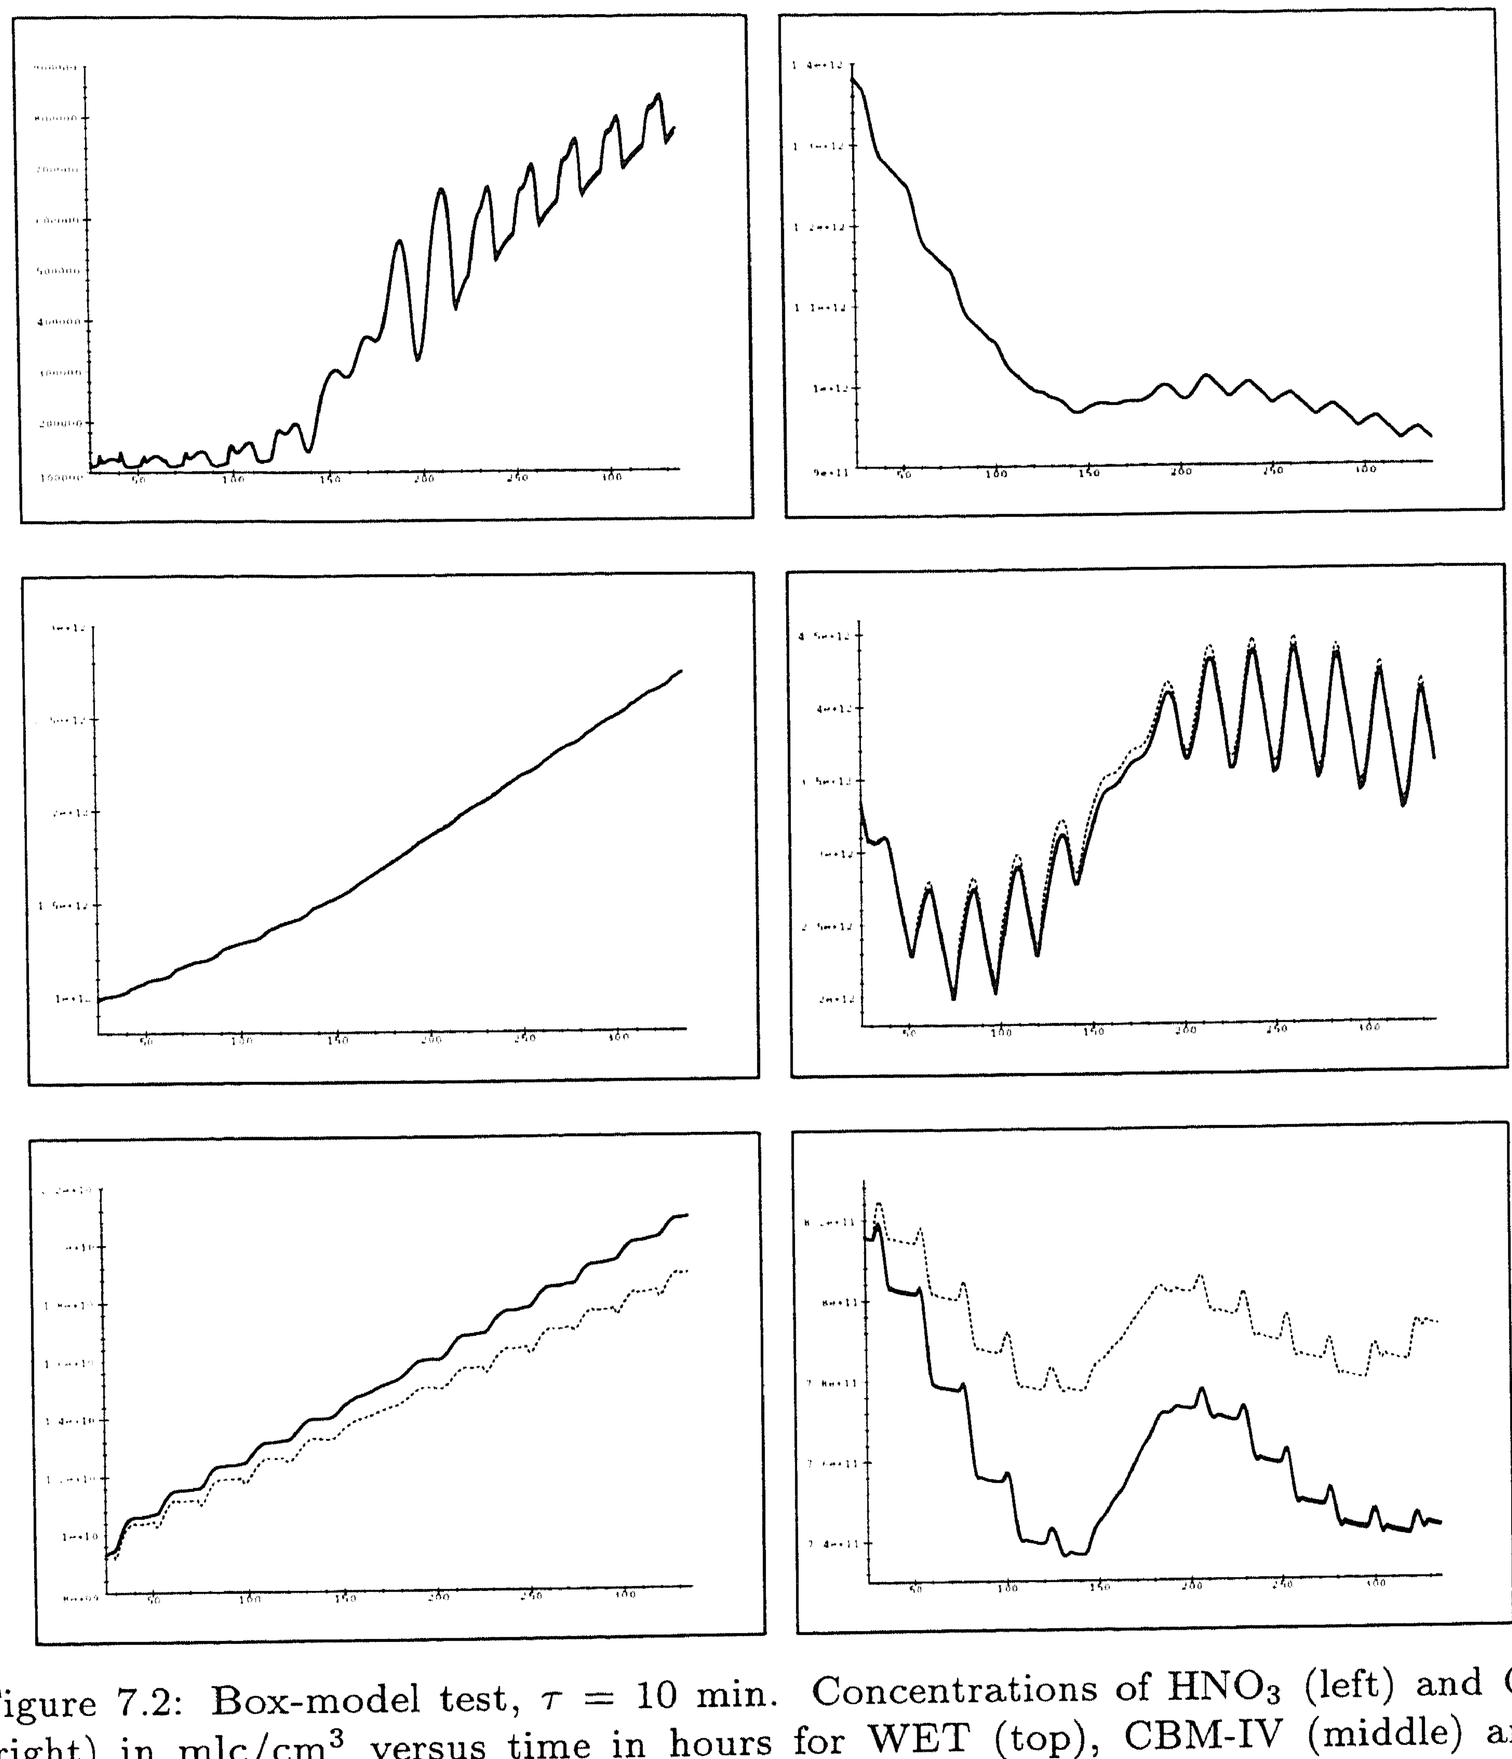

Figure 7.2: Box-model test, $\tau = 10$ min. Concentrations of HNO_3 (left) and O_3 (right) in mlc/cm^3 versus time in hours for WET (top), CBM-IV (middle) and RIVM (bottom). Thick solid line represents the reference solution, thin solid line the ROS2 solution and dotted line the TWOSTEP solution. TWOSTEP has not been applied to WET.

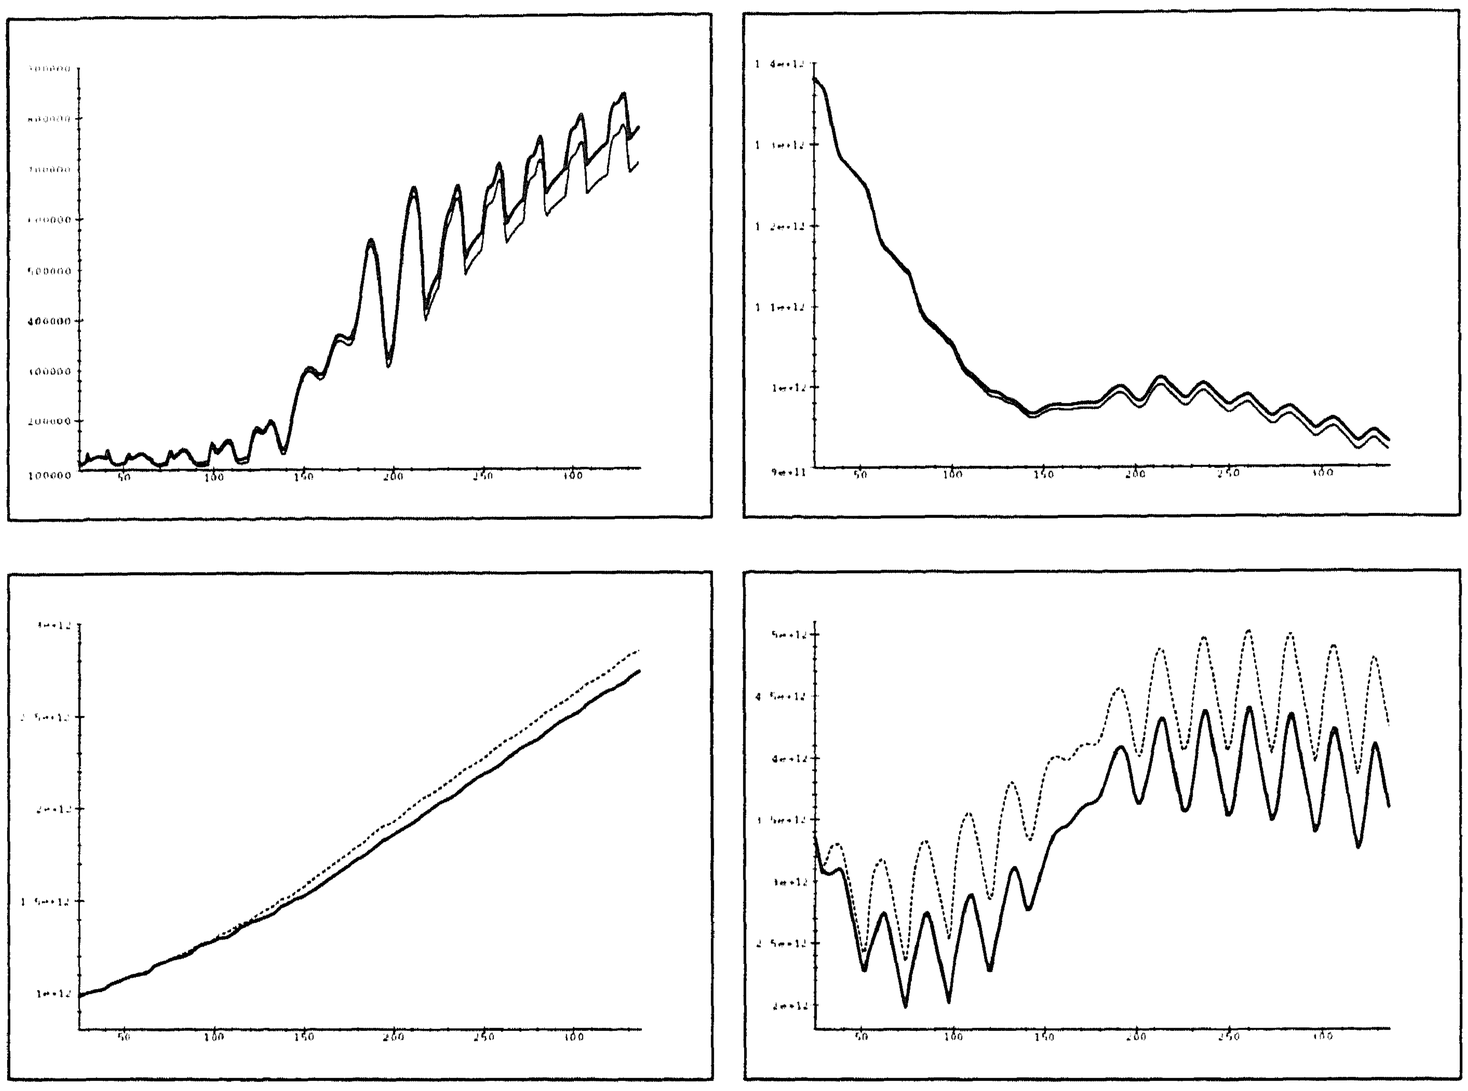

Figure 7.3: Box-model test, $\tau = 20$ min. Concentrations of HNO_3 (left) and O_3 (right) in mlc/cm^3 versus time in hours for WET (top), CBM-IV (middle) and RIVM (bottom). Thick solid line represents the reference solution, thin solid line the ROS2 solution and dotted line the TWOSTEP solution. TWOSTEP has not been applied to WET.

diffusion, here represented by F_A and F_D, respectively. Likewise, F_R stems from the chemical reactions, emissions and depositions. For any grid point the terms in F_R are simply the box-model expressions contained in f. We will discuss two second-order Strang-splittings, one treating vertical diffusion explicitly and one with an implicit vertical diffusion part.

7.4.1 Vertical diffusion explicit

For system (7.10) second-order Strang-splitting can be organized in many ways. We consider the combination

$$\begin{aligned} W_0 &= w_n, \\ W_1 &= W_0 + \tfrac{1}{4}\tau F_T(W_0) + \tfrac{1}{4}\tau F_T(W_0 + \tfrac{1}{2}\tau F_T(W_0)), \\ W_2 &= W_1 + \tfrac{3}{2}\tau k_1 + \tfrac{1}{2}\tau k_2, \\ W_3 &= W_2 + \tfrac{1}{4}\tau F_T(W_2) + \tfrac{1}{4}\tau F_T(W_2 + \tfrac{1}{2}\tau F_T(W_2)), \\ w_{n+1} &= W_3, \end{aligned} \tag{7.11}$$

where k_1 and k_2 are given by

$$(I - \gamma\tau A)\, k_1 = F_R(W_1), \quad (I - \gamma\tau A)\, k_2 = F_R(W_1 + \tau k_1) - 2k_1$$

and $A = F_R'(W_1)$. Hence the reaction part is treated with ROS2 and the transport part by the explicit trapezoidal rule. We note in passing that this rule derives from ROS2 by substitution of the zero matrix for A. By standard operator splitting we thus solve transport and chemistry in a sequential, symmetric manner such that the chemistry computation becomes completely decoupled from the transport. An advantage of standard splitting is that it amounts to chemistry box-model computations over the space grid and that it is easy to implement and memory efficient. A disadvantage is that the decoupling can result in stiff transients within any split step, as the decoupled transport changes the solution values for the chemistry integration. These stiff transients are a numerical artefact and may complicate the chemistry integration.

Stability

By using the explicit trapezoidal rule for the advective-diffusive transport, we tacitly assume that this does not lead to severe stability restrictions on the time step. In [HKvLV95] and [SVdZ+97a] stability of the explicit trapezoidal rule has been discussed for pure advection when using third-order upwind discretization with flux limiting. A CFL number of $\frac{2}{3}$ was shown to lead to a stable and positive advection computation. For practical purposes this is quite satisfactory. When also vertical diffusion is included, necessary is that

$$\tau \max \frac{4K}{(\Delta r)^2} \approx 1.0. \tag{7.12}$$

For our model (7.1) this condition allows sufficiently large step sizes. For example, substitution of the values for K and Δr given in Section 7.2, yields $\tau \approx 700^2/(4\times 30) \approx$ 4100 seconds. In case of much finer vertical meshes or much larger values for K, the explicit trapezoidal rule will no longer be efficient and must be replaced to obtain an implicit vertical diffusion computation.

1D Results

We first show numerical results for the 1D diffusion-reaction system

$$\frac{\partial c}{\partial t} = \frac{\partial}{\partial r}\left(\rho K \frac{\partial}{\partial r}\left(\frac{c}{\rho}\right)\right) + f(c), \quad 0 \le r \le r_H, \tag{7.13}$$

obtained from (7.1) for zero velocities. These results are of interest in their own as they enable a comparison between the 1D Strang-splitting scheme and ROS2 directly applied to the semi-discrete 1D problem $\dot{w} = F(w) \equiv F_D(w) + F_R(w)$ using the full Jacobian matrix. Needless to say that in 3D this is not feasible.

The same three chemistry models are used as in the box-model tests, now simulating the evolution of an air column along the trajectory of Figure 7.1. The initial values for c are chosen such that the mixing ratios c/ρ are independent of r, while the initial values at ground level are the same as for the box-models. Step sizes are chosen in the same way as in the box-model tests, i.e., $\tau = 10$ min and $\tau = 20$ min during day 2 up to day 14.

We found that for all three chemistry models the 1D ROS2 scheme and the 1D Strang-splitting scheme (7.11) are very close in accuracy. In fact, all computed solutions are in excellent agreement with the reference solution, being the semi-discrete solution at ground level computed in high time step accuracy. Figure 7.4 shows the evolution of O_3 and HNO_3 along the trajectory at ground level for the reference solution and the two numerical methods using $\tau = 20$ min. One can see that overall the agreement is almost up to plotting accuracy. This indicates that for the current 1D problem Strang-splitting has no adverse effect on accuracy whatsoever for step sizes smaller than or equal to 20 minutes.

It is stressed that when there are a large number of species, the use of the full banded Jacobian in ROS2 is not advocated in real practice, as this involves a considerable numerical algebra overhead (see also [VB96]). For WET, CBM-IV and RIVM the Strang-splitting computation is about a factor 19, 6 and 4 less expensive in CPU time.

3D Results

In 3D the Strang-splitting method (7.11) has been applied only for the chemical reaction set RIVM. The test is similar to the one carried out in [SVdZ+97a] on the coarsest reduced 64×32 horizontal space grid, yielding a total of 24840 grid cells in 3D. The same initial distribution was used. On this grid an accurate, semi-discrete

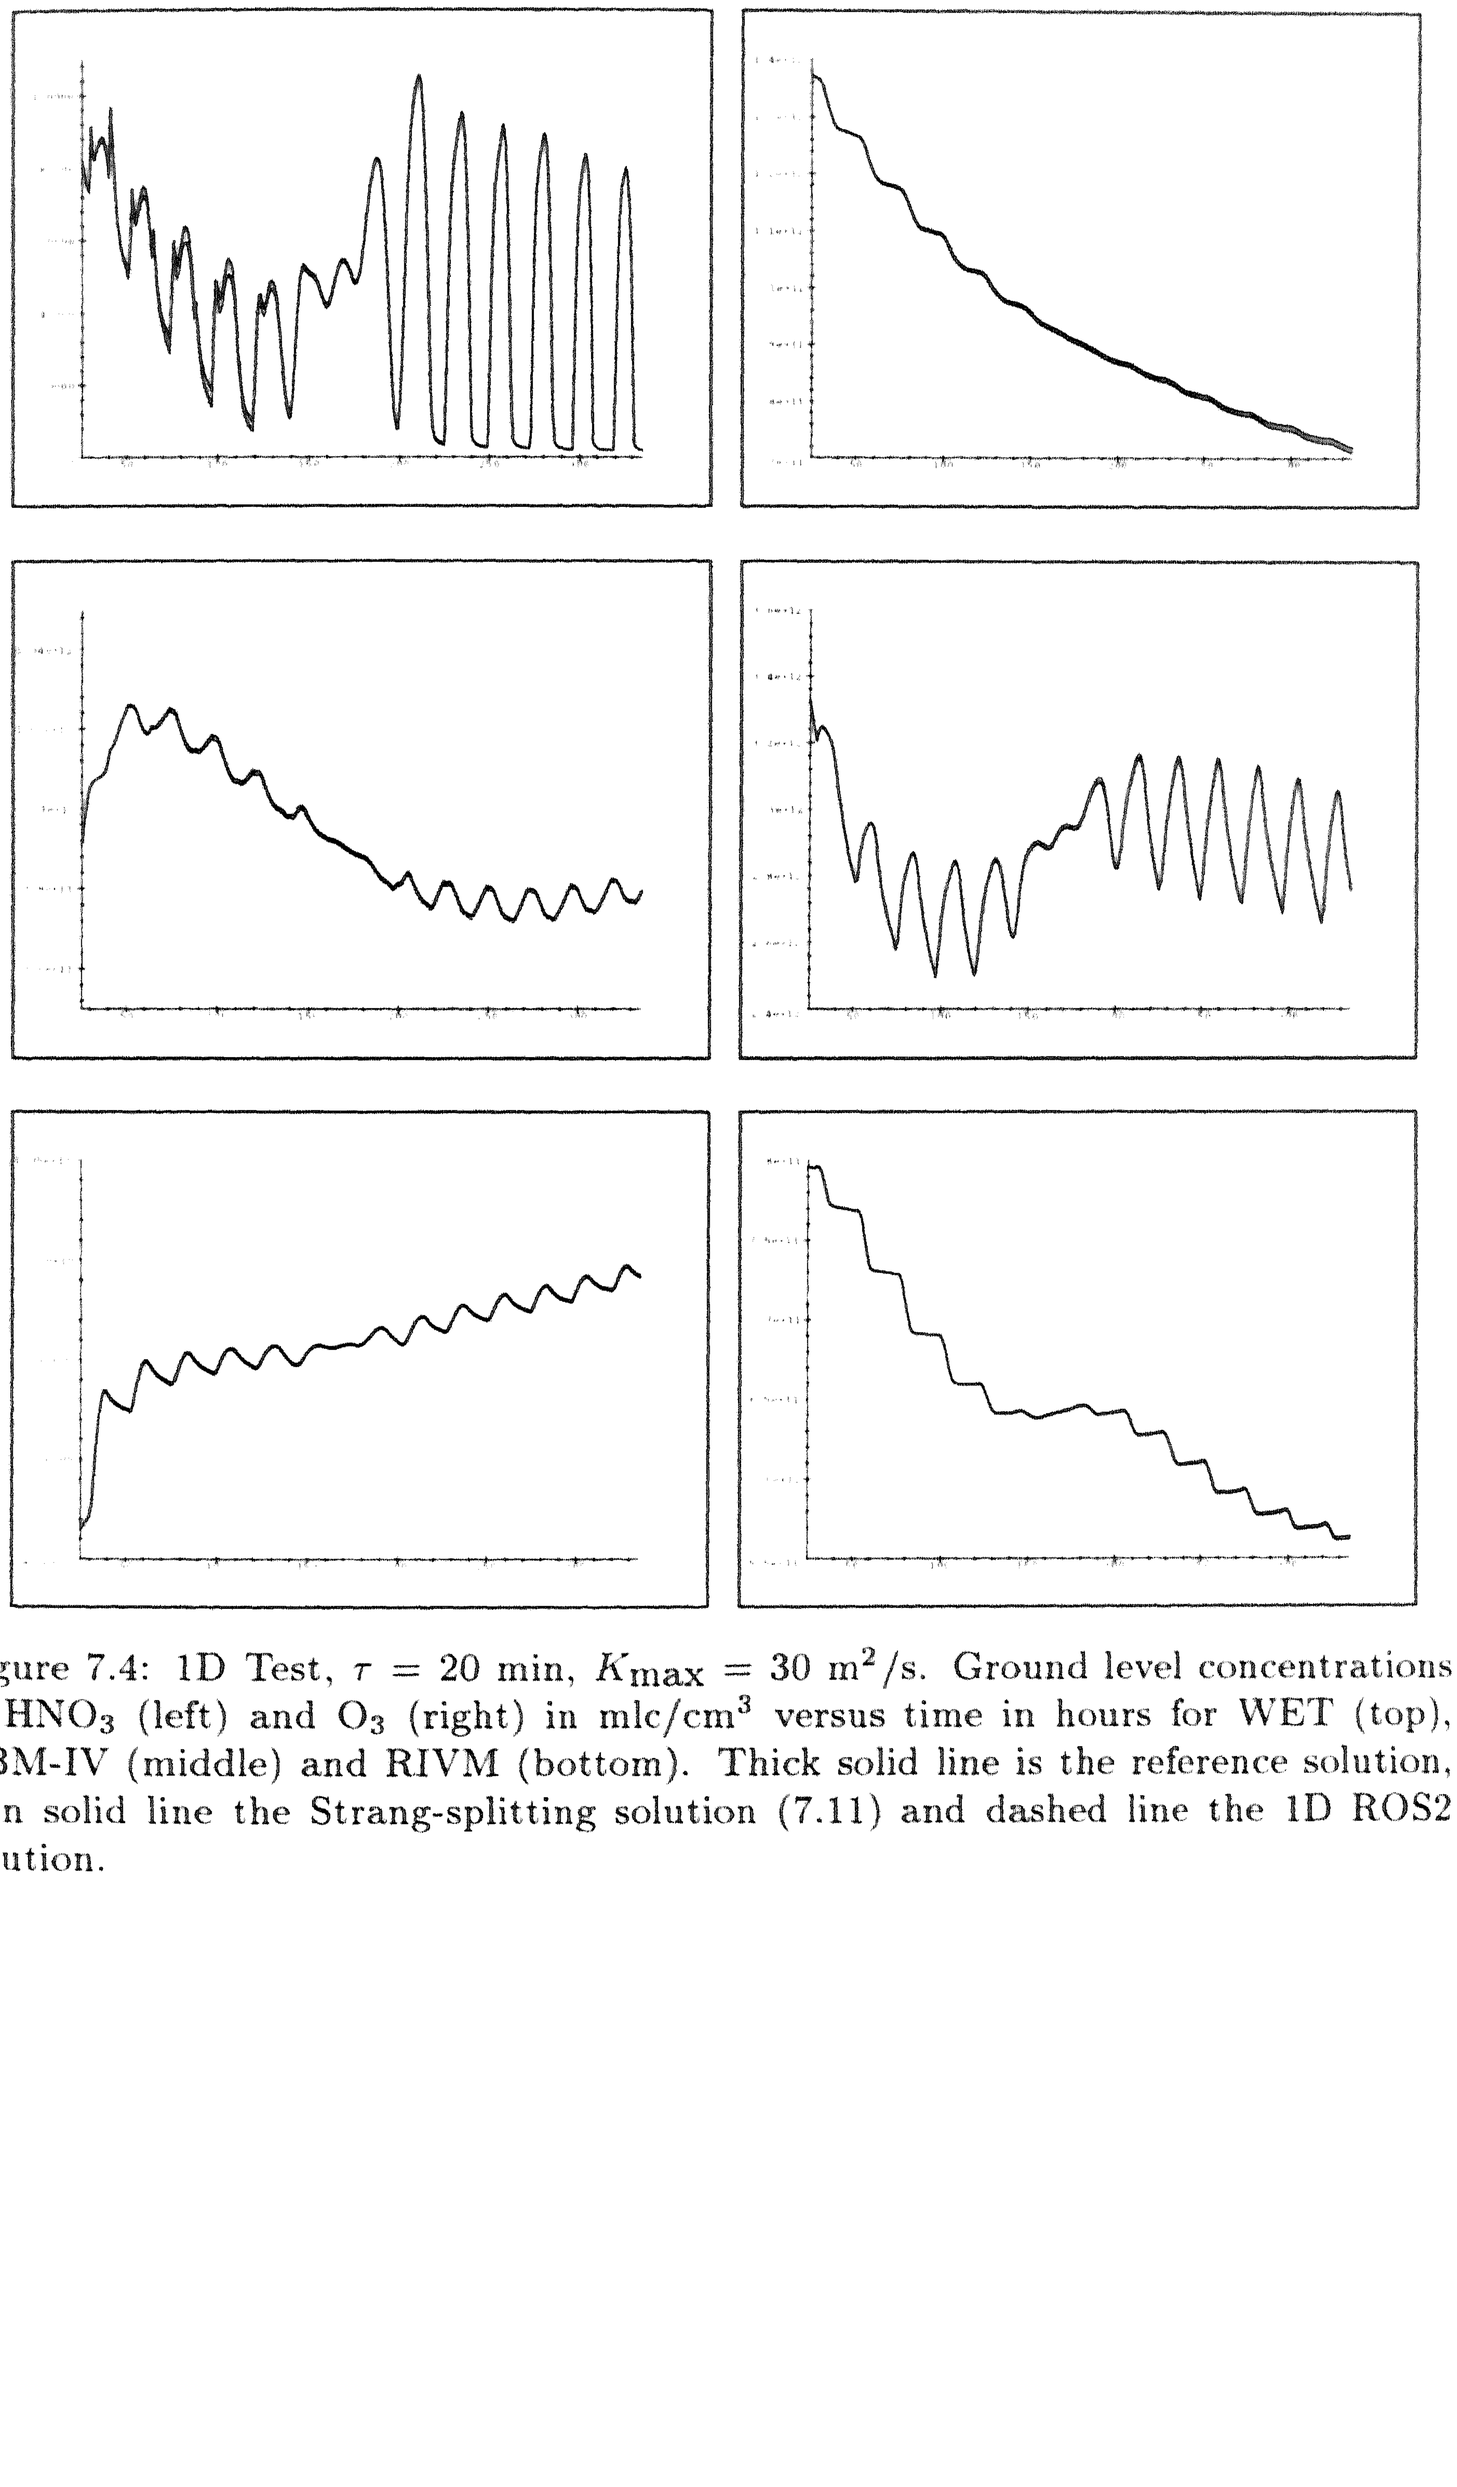

Figure 7.4: 1D Test, $\tau = 20$ min, $K_{\max} = 30$ m^2/s. Ground level concentrations of HNO_3 (left) and O_3 (right) in mlc/cm^3 versus time in hours for WET (top), CBM-IV (middle) and RIVM (bottom). Thick solid line is the reference solution, thin solid line the Strang-splitting solution (7.11) and dashed line the 1D ROS2 solution.

reference solution has been determined to furnish the initial values for the actual integration starting at day 2, as well as to assess the accuracy at the final time at day 14. We again successfully used fixed step sizes of 10, 20 min. These step sizes are allowed for the advection computation since on the chosen grid the critical step size for a stable advection computation is about 20 min. Note that in (7.11) the step size τ is halved in the transport steps.

The results clearly indicate that also in a 3D Strang-splitting code ROS2 is able to integrate the chemical kinetic equations with large step sizes. For an accuracy assessment we refer to Figure 7.5. This figure shows ground level profiles of O_3, HO_2NO_2 and HNO_3, plotted along the horizontal SW-NE grid diagonal of Figure 7.1 at the final time at day 14. For each species two profiles were plotted, the computed one and the reference profile. Only for HO_2NO_2 the errors are notable.

7.4.2 Vertical diffusion implicit

It is of numerical interest to examine a test case which requires an implicit vertical diffusion computation. For this purpose we have artificially increased the diffusion coefficient K in the prototype model by a factor 100 so that the maximum value has become 3000. The explicit trapezoidal rule is then no longer efficient for use in (7.11) and must be replaced. We have replaced it by a new scheme which also keeps the advective and diffusive transport coupled.

The new transport scheme is derived from ROS2 and exploits the fact that this Rosenbrock method is second-order consistent for any choice of the Jacobian approximation A. Specifically, we apply ROS2 to the transport problem $\dot{w} = F_T(w) \equiv F_A(w) + F_D(w)$ and choose $A = F_D'(w_n)$ to obtain

$$\begin{aligned} w_{n+1} &= w_n + \tfrac{3}{2}\tau k_1 + \tfrac{1}{2}\tau k_2, \\ (I - \gamma\tau F_D'(w_n))\, k_1 &= F_T(w_n), \\ (I - \gamma\tau F_D'(w_n))\, k_2 &= F_T(w_n + \tau k_1) - 2k_1. \end{aligned} \tag{7.14}$$

Hence advection is still treated explicitly while, owing to the 1D nature of F_D', the computation of the stage vectors k_1, k_2 now requires the solution of tridiagonal linear systems, one for each species and each horizontal grid point. With respect to CPU this scheme is therefore almost as cheap per step as the explicit trapezoidal rule. Replacing the explicit trapezoidal rule in (7.11) by (7.14) gives the second-order Strang-splitting scheme

$$\begin{aligned} W_0 &= w_n, \\ W_1 &= W_0 + \tfrac{3}{4}\tau k_1^{(1)} + \tfrac{1}{4}\tau k_2^{(1)}, \\ W_2 &= W_1 + \tfrac{3}{2}\tau k_1^{(2)} + \tfrac{1}{2}\tau k_2^{(2)}, \\ W_3 &= W_2 + \tfrac{3}{4}\tau k_1^{(3)} + \tfrac{1}{4}\tau k_2^{(3)}, \\ w_{n+1} &= W_3, \end{aligned} \tag{7.15}$$

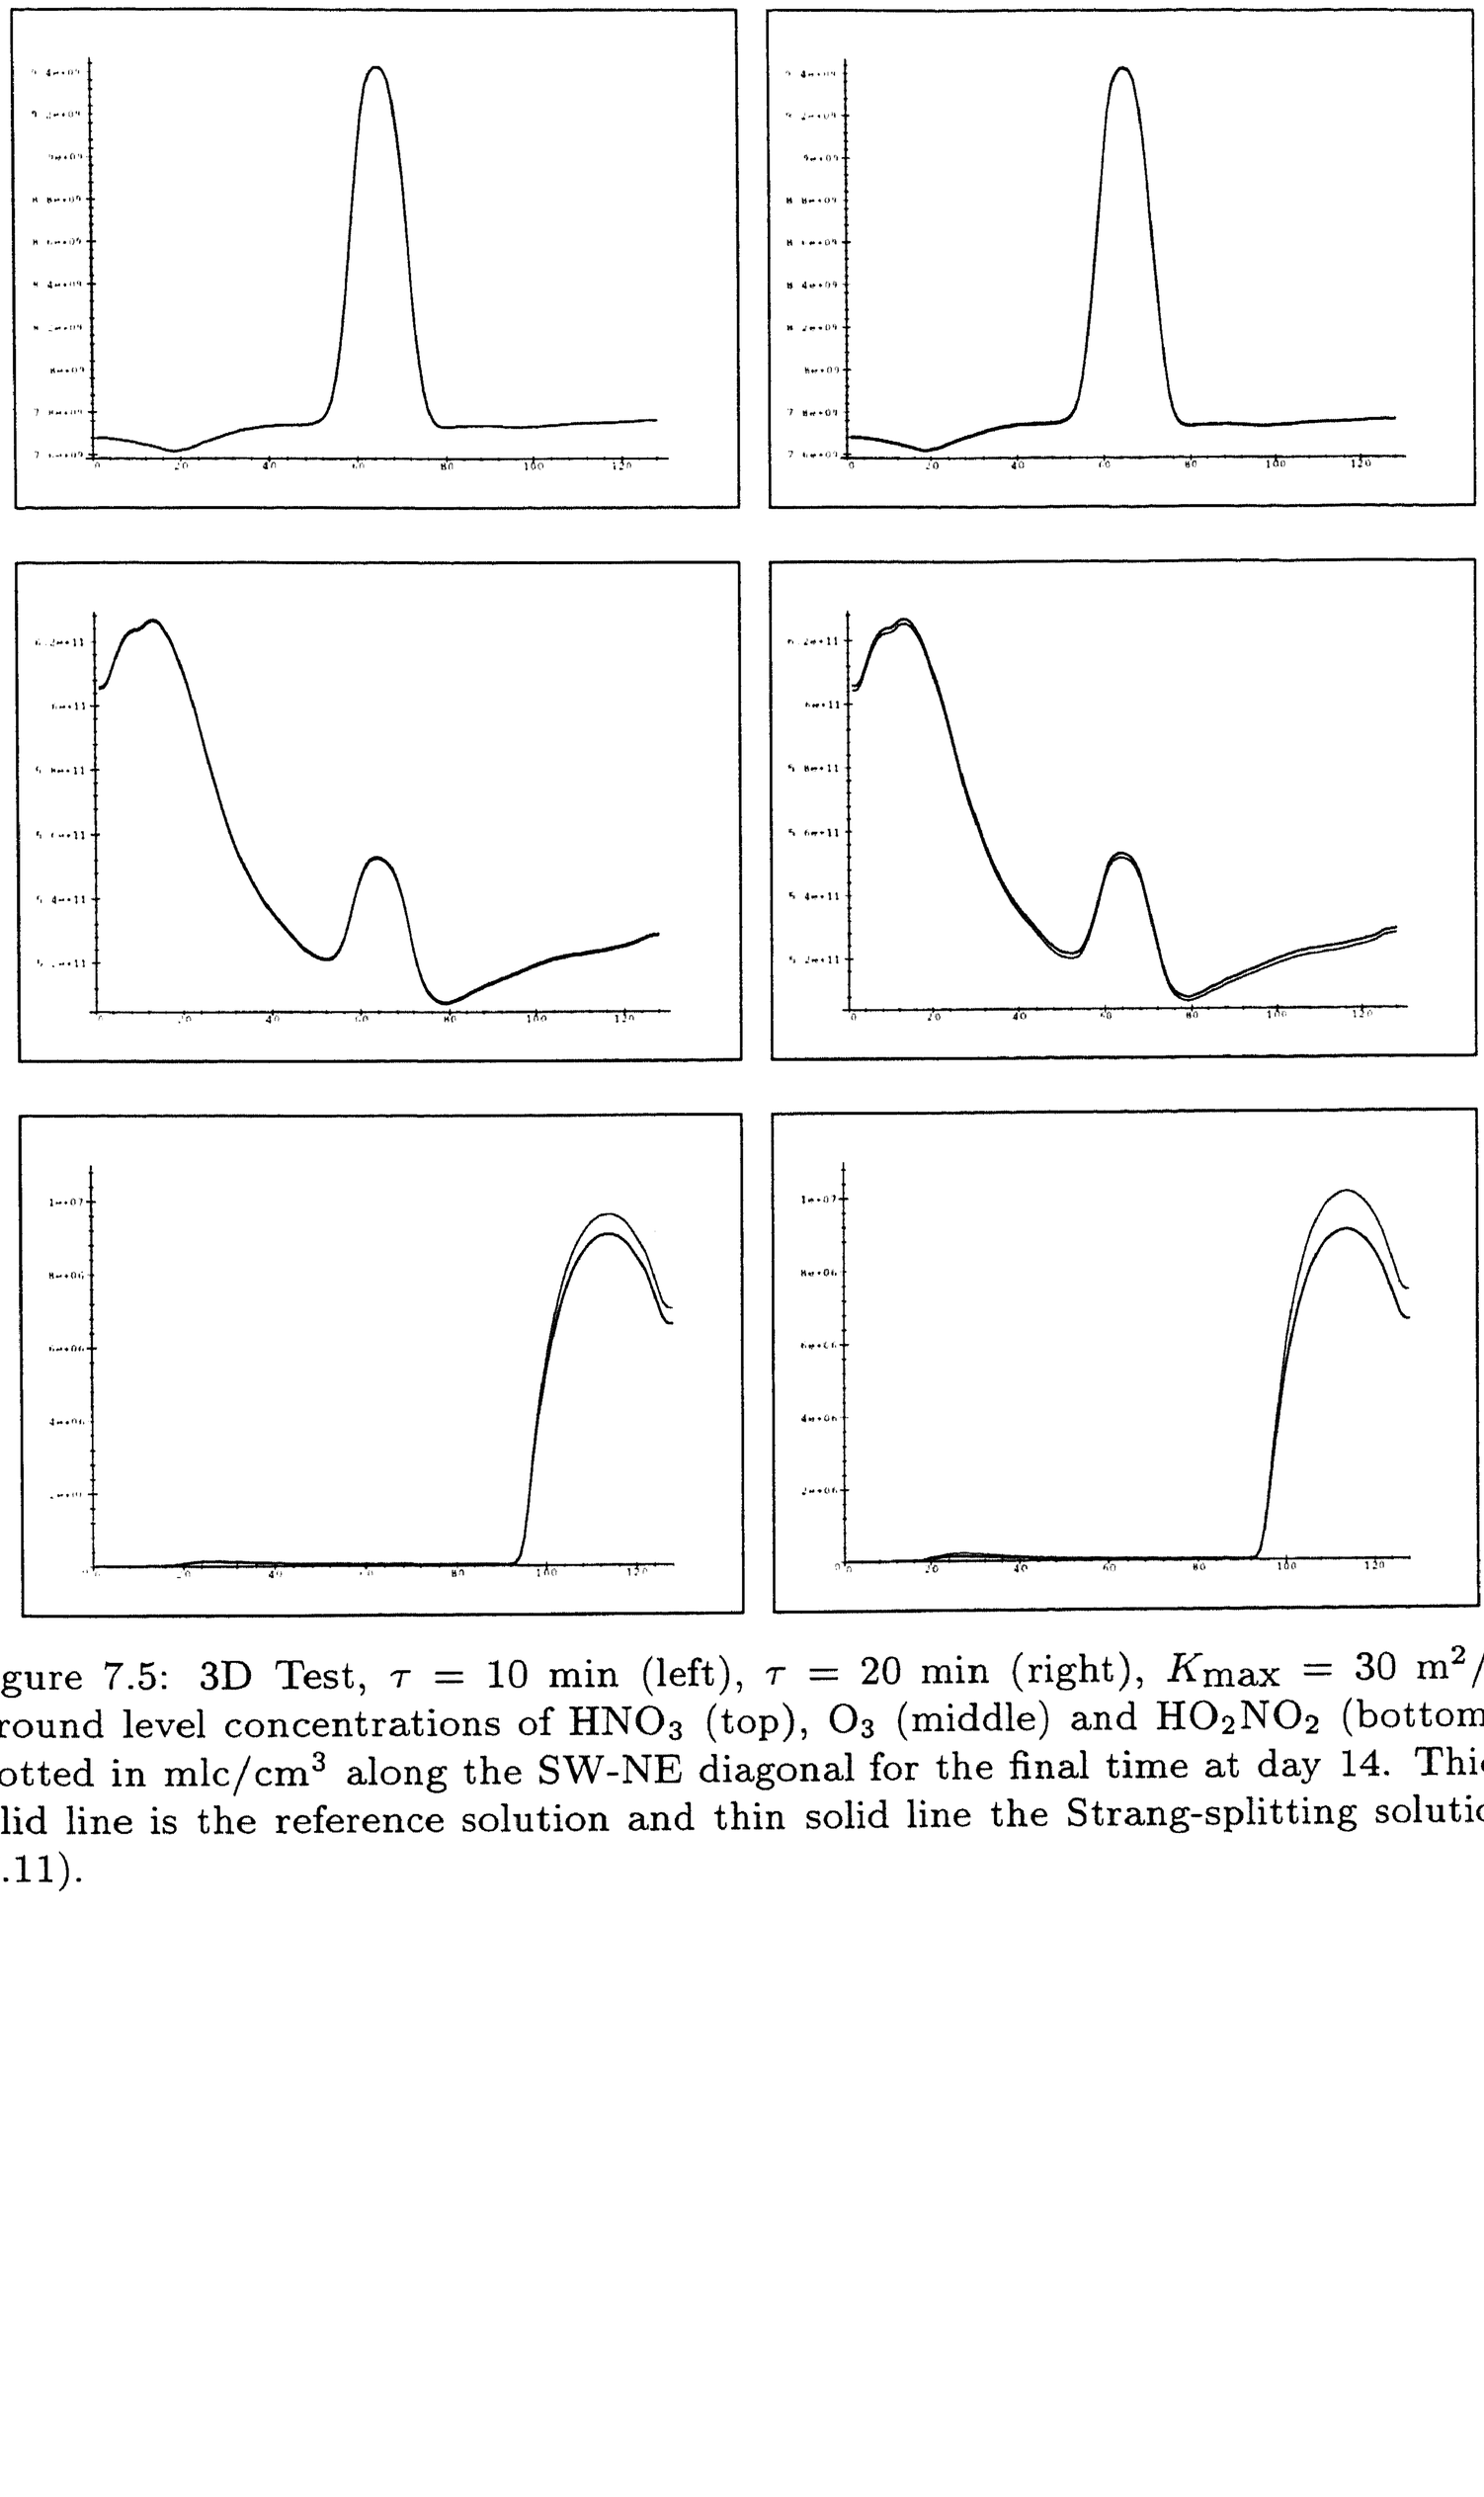

Figure 7.5: 3D Test, $\tau = 10$ min (left), $\tau = 20$ min (right), $K_{\max} = 30$ m^2/s. Ground level concentrations of HNO_3 (top), O_3 (middle) and HO_2NO_2 (bottom), plotted in mlc/cm^3 along the SW-NE diagonal for the final time at day 14. Thick solid line is the reference solution and thin solid line the Strang-splitting solution (7.11).

where W_1 and W_3 are computed by (7.14) and W_2 is obtained in the same way as in (7.11).

Stability

Using an approximation for the Jacobian in the Rosenbrock method of course affects the stability. Following the standard way of reasoning, for the advection-diffusion scheme (7.14) we examine the stability for the scalar, linear test model

$$\dot{w} = \lambda_A w + \lambda_D w. \tag{7.16}$$

Through an eigenvector-eigenvalue decomposition this scalar model is derived from the constant coefficient linear system

$$\dot{w} = J_A w + J_D w,$$

where J_A, J_D represent 'frozen' Jacobians $F'_A(w_n), F'_D(w_n)$. Hence, λ_A represents a complex-valued eigenvalue of the advection Jacobian $F'_A(w_n)$. Likewise, λ_D represents a real, nonpositive eigenvalue of the diffusion Jacobian $F'_D(w_n)$. The eigenvector-eigenvalue decomposition is valid if the matrices share the same eigensystem in which case it can also be carried through for the Rosenbrock approximations. Assuming spatial discretization on a uniform grid, this decomposition holds for problems with constant coefficients and periodic boundary conditions.

Denote $z_A = \tau\lambda_A$ and $z_D = \tau\lambda_D$. The stability function of (7.14) then is

$$R(z_A, z_D) = \frac{(1 + z_A + \frac{1}{2}z_A^2) + (1 - 2\gamma)(1 + z_A)z_D}{(1 - \gamma z_D)^2}, \tag{7.17}$$

which reduces to the stability function $R(z_A, 0) = 1 + z_A + \frac{1}{2}z_A^2$ of the explicit trapezoidal rule for $z_D = 0$ and to (7.5) for $z_A = 0$. Because we assume that advection can be computed explicitly, we tacitly assume that $z_A = \mathcal{O}(1)$. As a consequence, $R(z_A, z_D) \to 0$ for $z_D \to -\infty$. Thus the damping at infinity property of the original L-stable stability function (7.5) is maintained for the large and negative diffusion eigenvalues.

Lemma 1. $|R(z_A, z_D)| \leq 1$ whenever $z_D \leq 0$ is real and $|1 + z_A + \frac{1}{2}z_A^2| \leq 1$.

Proof. Denote $\theta = (2\gamma - 1)/\gamma = \pm\sqrt{2}$. Using the notation $\tilde{z} = \gamma z_D$, $\alpha = 1 + z_A + \frac{1}{2}z_A^2$ and $\beta = \theta(1 + z_A)$, we can write

$$R(z_A, z_D) = \frac{\alpha - \beta\tilde{z}}{(1 - \tilde{z})^2}.$$

Suppose $\tilde{z} = -t \leq 0$. Then

$$|R(z_A, z_D)|^2 = \frac{|\alpha|^2 + \eta t + |\beta|^2 t^2}{(1 + t)^4}, \qquad \eta = \alpha\bar{\beta} + \bar{\alpha}\beta = 2\,\mathrm{Re}(\alpha\bar{\beta}).$$

So, if $|\alpha| = 1$, we have $|R(z_A, z_D)| \leq 1$ for all $t \geq 0$ iff

$$\eta + |\beta^2| t \leq 4 + 6t + 4t^2 + t^3 \qquad \text{for all } t \geq 0.$$

Below it will be shown that $|\alpha| = 1$ implies

$$\eta \leq \tfrac{3}{2}\sqrt{6} < 4 \qquad \text{and} \qquad |\beta|^2 \leq 6, \tag{7.18}$$

which proves the lemma.

To verify the inequalities in (7.18), let $\zeta = 1 + z_A$. Then $\alpha = \frac{1}{2}(1 + \zeta^2)$, $\beta = \theta\zeta$ and

$$\eta = \theta(1 + |\zeta|^2) \operatorname{Re} \zeta.$$

We consider values of z_A such that $|\alpha| = 1$. Hence $|1 + \zeta^2| = 2$ and it easily follows that $|\beta|^2 = 2|\zeta|^2 \leq 6$. Further, any ζ satisfying $|1 + \zeta^2| = 2$ can be parameterized as

$$\zeta = \cos\phi + i\sqrt{\kappa}\,\sin\phi, \qquad \kappa = 1 + 2|\sin\phi|^{-1}.$$

By some calculations we thus arrive at $\eta = 2\theta(1 + |\sin\phi|)\cos\phi$ and the maximum value for η is found to be $\frac{3}{2}\sqrt{6}$. □

This lemma proves unconditional stability for all real, nonpositive z_D as long as the complex number z_A lies in the stability region of the explicit trapezoidal rule. Hence, for the linear model problem the critical step size for stability is equal to the critical step size for advection. The advection-diffusion solver (7.14) therefore is of interest in its own and can also prove useful in other applications.

1D Results

We have repeated the 1D test for problem (7.13) with the 100 times larger diffusion coefficient K, again comparing the 1D Strang-splitting scheme (7.15) to the 1D ROS2 scheme for $\tau = 10$, 20 min. Figure 7.6 shows the O_3 and HNO_3 profiles at ground level, in the same manner as in the previous 1D Figure 7.4. We see that for all three chemistry models the accuracy of the splitting scheme is still good but it has become lower, especially for $\tau = 20$ min. This is due to the much larger value for K. The 1D ROS2 scheme performs equally well and accurate for WET and RIVM. However, it has become unstable for the chemistry model CBM-IV. This happens for both step sizes. This instability is odd in view of the fact that splitting gives a fairly accurate, stable result. For the CBM-IV model the 1D ROS2 plot has been omitted.

3D Results

Figure 7.7 shows results of the 3D test using the 100 times larger diffusion coefficient K. The new Strang-splitting scheme (7.15) can be seen to solve the problem well for both step sizes of 10 and 20 min (In Figure 7.7 also results are included for a different

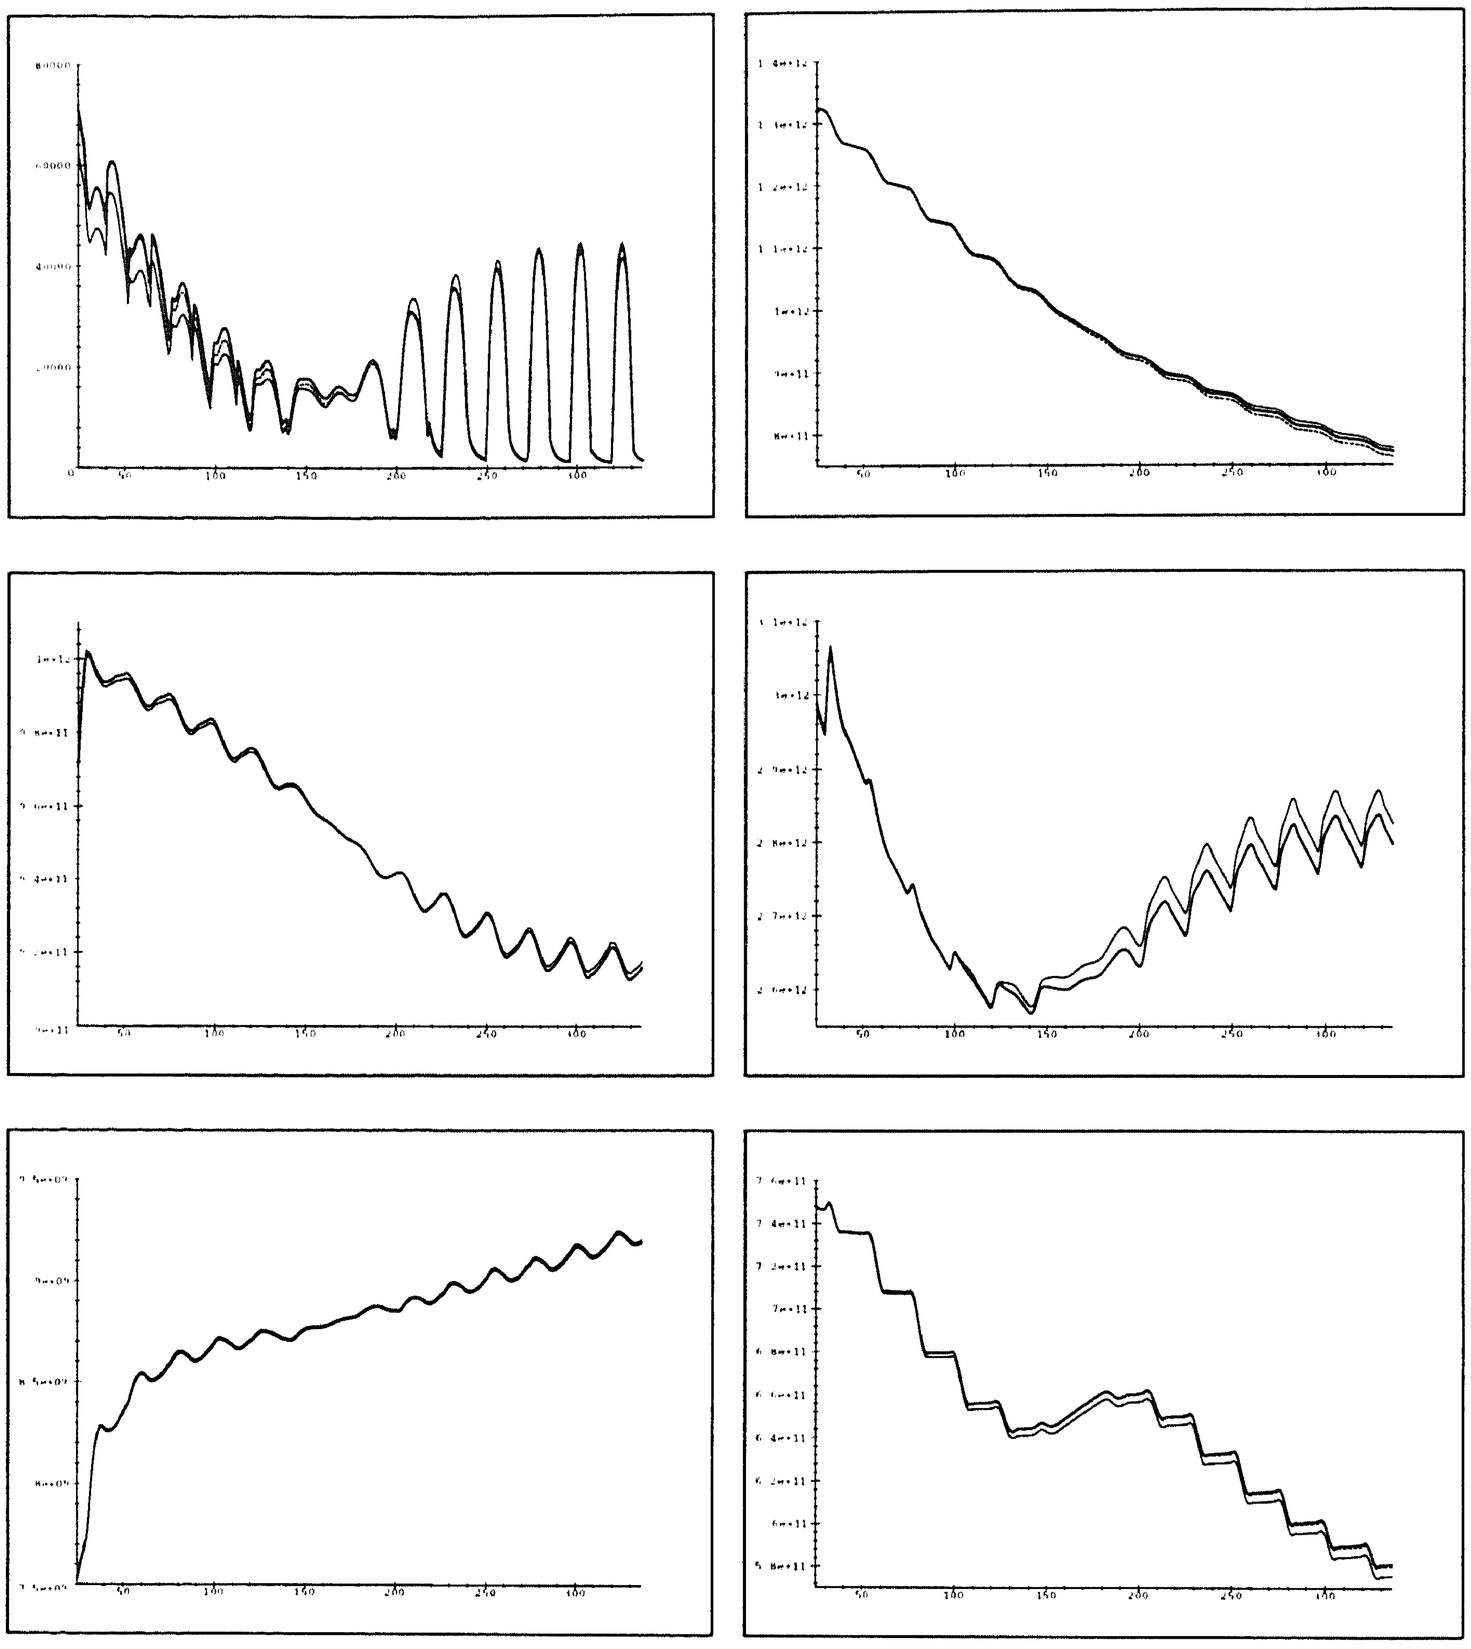

Figure 7.6: 1D Test, $\tau = 20$ min, $K_{\max} = 3000$ m^2/s. Ground level concentrations of HNO_3 (left) and O_3 (right) in mlc/cm^3 versus time in hours for WET (top), CBM-IV (middle) and RIVM (bottom). Thick solid line is the reference solution, thin solid line the Strang-splitting solution (7.15) and dashed line the 1D ROS2 solution.

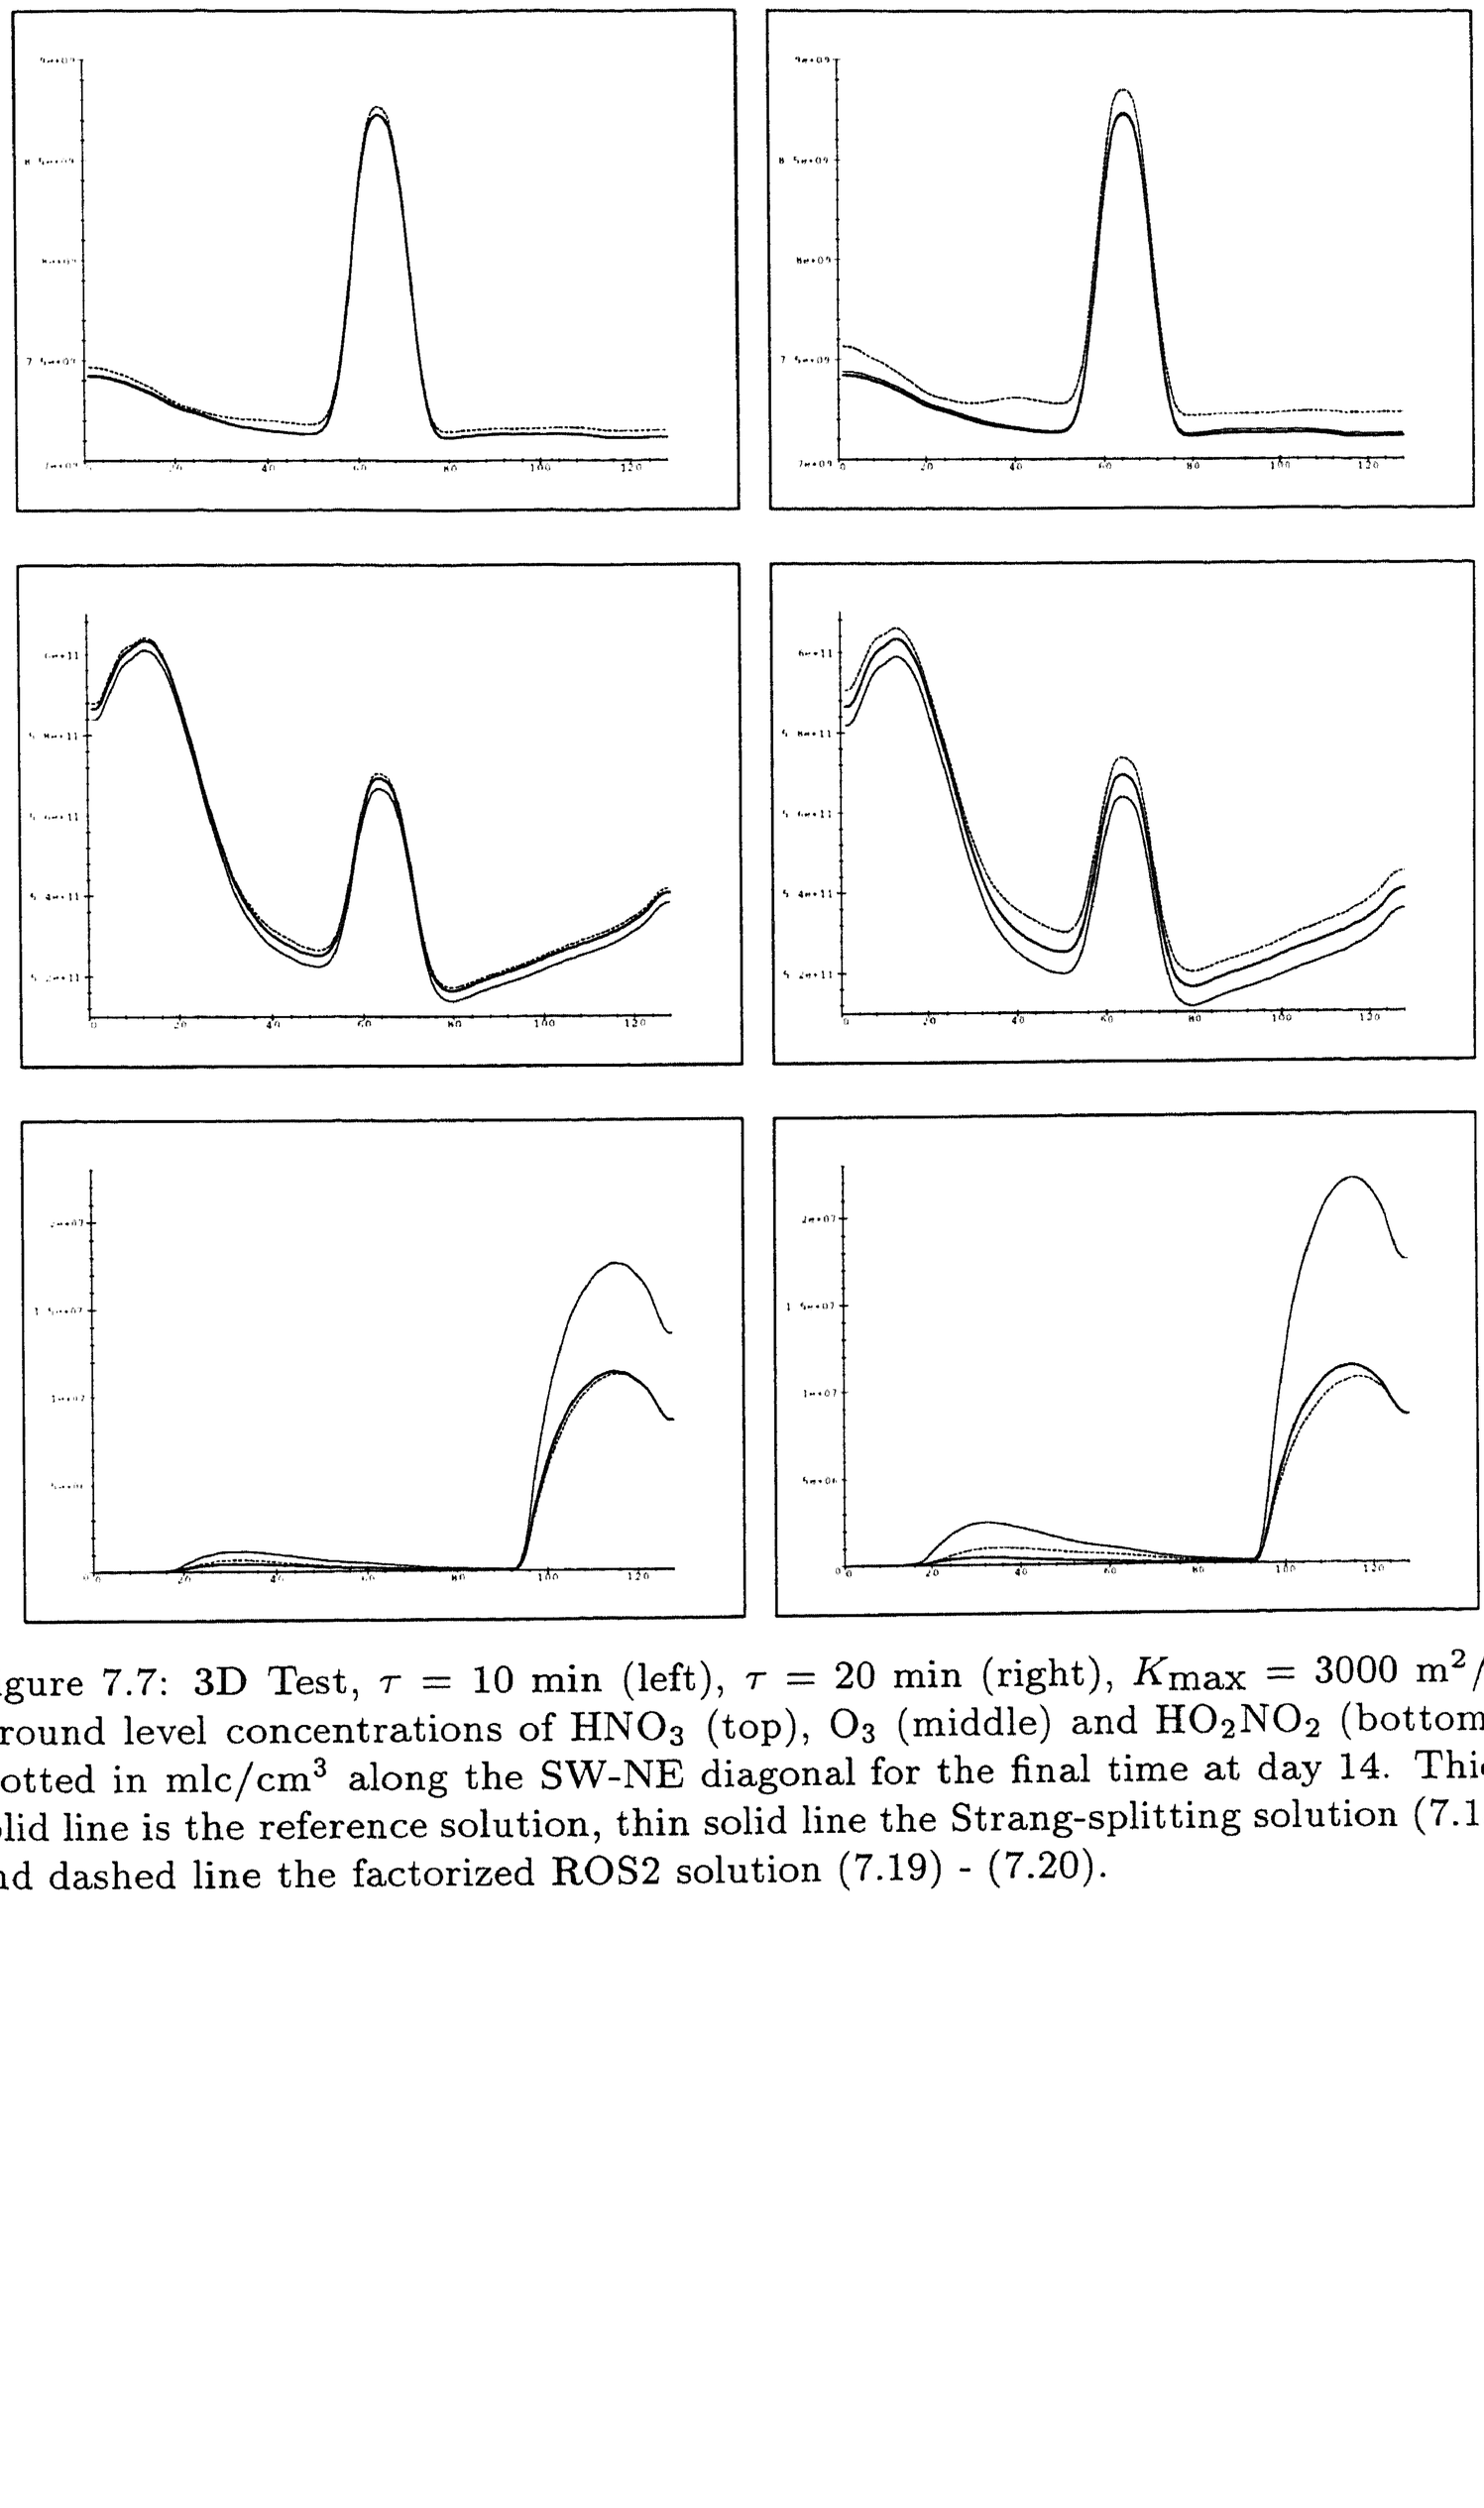

Figure 7.7: 3D Test, $\tau = 10$ min (left), $\tau = 20$ min (right), $K_{\max} = 3000$ m^2/s. Ground level concentrations of HNO_3 (top), O_3 (middle) and HO_2NO_2 (bottom), plotted in mlc/cm^3 along the SW-NE diagonal for the final time at day 14. Thick solid line is the reference solution, thin solid line the Strang-splitting solution (7.15) and dashed line the factorized ROS2 solution (7.19) - (7.20).

method that will be discussed in Section 7.5). We do encounter larger errors though compared to those of the previous 3D figure. The error is very large for HO_2NO_2, but for the most important species O_3 it remains within acceptable bounds.

It is emphasized that the larger errors are due to the splitting and that the transport and chemistry schemes themselves do hardly contribute to the observed errors. In other words, replacing these two schemes by the exact solution operators for the transport problem $\dot{w} = F_T(w)$ and the chemistry problem $\dot{w} = F_R(w)$, within the framework of Strang-splitting, will not annihilate the error.

The splitting error has also been observed in related 3D tests carried out in [SVdZ+97a], where a completely different chemical integrator based on TWOSTEP has been used. Because we artificially increased the expression for K from [War88] by a factor of 100, we must admit that we are not sure whether this 3D test is really meaningful for the actual practice of computational air quality modeling. On the other hand, it is a sound numerical test which has confirmed the accuracy and robustness of the Rosenbrock method ROS2 as an efficient chemical integrator within Strang-splitting.

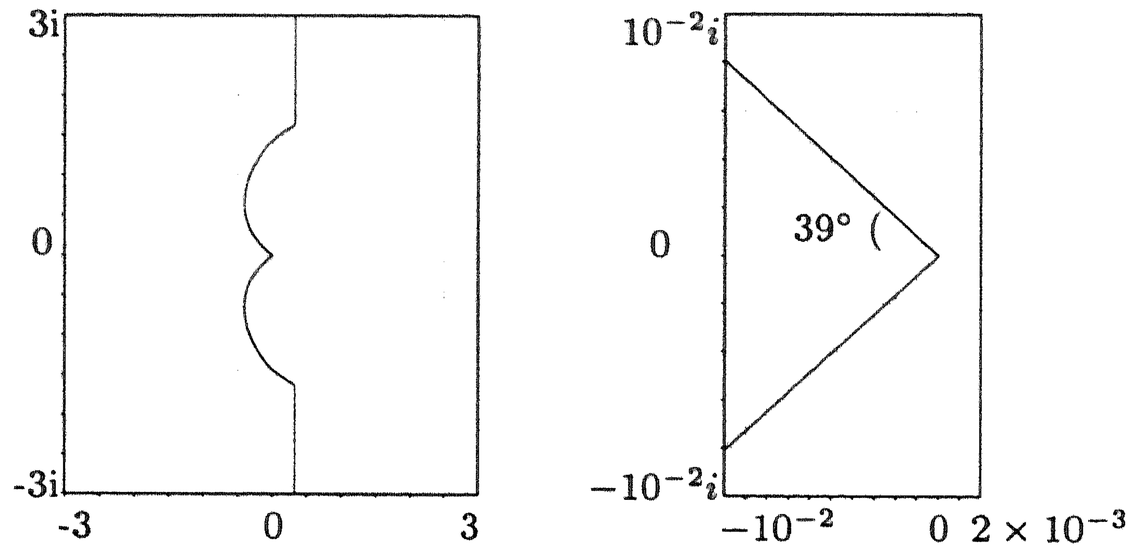

Figure 7.8: At the left the z_R-stability boundary for the factorized stability function (7.22). At the right the magnified subregion near the origin.

7.5 ROS2 applied with approximate factorization

The main idea behind operator splitting is to avoid the complications of solving the huge systems of linear and nonlinear algebraic equations encountered when applying implicit or linearly-implicit time integration methods. Somehow related to splitting is the notion of approximate factorization, where a form of splitting is performed

at the numerical algebra level rather than at the operator level. See [AB96, BW76, WB79, vdHSK96] for examples of approximate factorization. By splitting at the numerical algebra level the operator splitting error is avoided. In this section we will therefore briefly examine whether for our 3D photochemical dispersion problem the Rosenbrock method ROS2 applied with a certain approximate factorization can offer a viable alternative to Strang-splitting for problems with large vertical diffusion.

As before we consider the ODE system

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

for which ROS2 reads

$$\begin{aligned} w_{n+1} &= w_n + \tfrac{3}{2}\tau k_1 + \tfrac{1}{2}\tau k_2, \\ (I - \gamma\tau A)\, k_1 &= F(w_n), \\ (I - \gamma\tau A)\, k_2 &= F(w_n + \tau k_1) - 2k_1. \end{aligned} \tag{7.19}$$

Approximate factorization is applied by defining A such that $I - \gamma\tau A$ is factorized into

$$I - \gamma\tau A = (I - \gamma\tau F_D'(w_n))\,(I - \gamma\tau F_R'(w_n)). \tag{7.20}$$

The computation of the two stage vectors k_1, k_2 then amounts to a normal F-evaluation and two sequential linear system solutions, one for the vertical diffusion and one for the chemistry. Owing to the box-model structure of F_R' and the tridiagonal structure of F_D' (observe the resemblance with (7.14)), these linear systems solutions can be carried out efficiently. The normal F-evaluation avoids the decoupling of operator splitting and the two linear system solutions are introduced for stability reasons. The order in the factorization is important. If we reverse the order, i.e. first solve for the chemistry, the numerical performance will decrease.

7.5.1 Stability

Similar as for (7.14), a separate stability investigation of (7.19) is required. For that purpose we again consider the standard scalar test model, now in the form

$$\dot{w} = \lambda_A\, w + \lambda_D\, w + \lambda_R\, w. \tag{7.21}$$

Denote $z_A = \tau\lambda_A$, etc. and $z = z_A + z_D + z_R$. The stability function of the factorized ROS2 scheme then can be written as

$$R(z_A, z_D, z_R) = 1 + \frac{2z}{(1 - \gamma z_D)\,(1 - \gamma z_R)} + \frac{\frac{1}{2}z^2 - z}{(1 - \gamma z_D)^2\,(1 - \gamma z_R)^2}. \tag{7.22}$$

Observe that due to the factorization there is no longer damping at infinity. For $z_D = z_R = 0$ the stability function of the explicit trapezoidal rule is recovered.

We are interested in stability whenever $z_D < 0$ is real and the explicit trapezoidal rule stability function satisfies $|1 + z_A + \frac{1}{2}z_A^2| \leq 1$. This leads to a stability region for z_R which we have determined numerically, see Figure 7.8. The figure reveals $A(\alpha)$-stability for z_R. Recall that a method is said to be $A(\alpha)$-stable if the sector

$$\{z : |\pi - \arg(z)| < \alpha, \ \ \mathrm{Re}(z) < 0\}$$

lies in the stability region. A careful inspection near the origin revealed that the angle α is very close to 39°. This angle probably is sufficiently large since eigenvalues with a large imaginary part do not seem to occur in atmospheric chemistry models. With a weaker condition on z_A a larger angle will be found. In conclusion, with respect to stability for the linear model problem, the critical step size is equal to that of the explicit trapezoidal rule for the advection computation. This means that with respect to model problem stability, the factorized ROS2 scheme has the same stability characteristics as the two previous Strang-splitting schemes.

7.5.2 3D Results

We have compared the factorized ROS2 scheme (7.19) - (7.20) to the Strang-splitting scheme (7.15) by repeating the 3D prototype model test with the 100 times larger diffusion coefficient K. Results are shown in Figure 7.7 in the same way as for the Strang-splitting scheme. As we anticipated, the huge error in HO_2NO_2 is absent now. On the other hand, for O_3 the errors are comparable and for HNO_3 the Strang-splitting solution is even slightly more accurate. Hence the factorized ROS2 scheme offers an improvement, but less than expected. In this respect it is emphasized that factorization as in (7.20) also introduces errors, as splitting does. By factorization we do use an approximation to the true Jacobian matrix $F'(w_n)$. Approximating $F'(w_n)$ in a Rosenbrock method leads to terms in the local truncation error different from the elementary differentials. The contribution of such new terms to the local truncation error is hard to predict, but it is most likely that they will increase the local error. With regard to CPU time the two schemes are comparable but Strang-splitting is somewhat more convenient for economical memory use.

7.6 Source-splitting

During our investigations we have examined another alternative for Strang-splitting, which we call source-splitting. Source-splitting is advocated in [Kes95, KW93, Pro97, Sun96]. The underlying idea is to treat the transport term as a constant source within the chemistry integration, so that a change of the solution values as happens in Strang-splitting is avoided. Similar as for Strang-splitting, the idea can be implemented in different ways. Adopting it for ROS2 as chemistry scheme and the explicit trapezoidal rule as transport scheme for the system $\dot{w} = F_T(w) + F_R(w)$, yields

$$w_{n+1} \quad = \quad w_n + \tfrac{3}{2}\tau k_1 + \tfrac{1}{2}\tau k_2, \tag{7.23}$$

$$\begin{aligned}(I-\gamma\tau A)\,k_1 &= F_R(w_n)+\tilde{F}_T,\\ (I-\gamma\tau A)\,k_2 &= F_R(w_n+\tau k_1)+\tilde{F}_T-2k_1,\end{aligned}$$

where $A = F_R'(w_n)$ and the source term $\tilde{F}_T$ is defined by

$$\tilde{F}_T = \tfrac{1}{2}F_T(w_n) + \tfrac{1}{2}F_T(w_n + \tau F_T(w_n)). \tag{7.24}$$

This scheme shares the implementational advantages of the related splitting scheme (7.11), but it is only of first-order.

Treating transport as a source term requires a separate stability investigation. The stability function for the test model $\dot{w} = \lambda_T w + \lambda_R w$ reads

$$R(z_T, z_R) = \frac{(1 + z_T + \frac{1}{2}z_T^2) + (1-2\gamma)z_R + \frac{1}{2}(1-4\gamma)z_R(z_T + \frac{1}{2}z_T^2)}{(1-\gamma z_R)^2}. \tag{7.25}$$

We have found that with respect to stability the method is less stable than (7.11). Below we prove that $|R(z_T, z_R)| \le 1$ whenever $z_R \le 0$ is real and $|1 + z_T + \frac{1}{2}z_T^2| \le 1$, but that no positive angle α exists for which this stability function is $A(\alpha)$-stable with respect to z_R uniformly for z_T in the stability region of the explicit trapezoidal rule. In other words, unlike the Strang-splitting method (7.11), the source-splitting method does not simultaneously maintain the stability provided by the explicit trapezoidal rule for advection-diffusion and the L-stability of ROS2 for the chemistry.

Lemma 2. $|R(z_T, z_R)| \le 1$ whenever $z_R \le 0$ is real and $|1 + z_T + \frac{1}{2}z_T^2| \le 1$.

Proof. Put $1 + z_T + \frac{1}{2}z_T^2 = e^{i\phi}$. Then

$$\begin{aligned}R(z_T, z_R) &= \frac{e^{i\phi} + (e^{i\phi}-1)(1-4\gamma)\frac{1}{2}z_R + (1-2\gamma)z_R}{(1-\gamma z_R)^2} \\ &= \frac{e^{i\phi}(1 + (1-4\gamma)\frac{1}{2}z_R) + \frac{1}{2}z_R}{(1-\gamma z_R)^2}.\end{aligned} \tag{7.26}$$

Suppose $z_R = -t, t \ge 0$ and write $\alpha = 1 - \frac{1}{2}(1-4\gamma)t$. Then

$$|R(z_T, z_R)|^2 = \frac{(\alpha\cos(\phi) - \frac{1}{2}t)^2 + \alpha^2\sin^2(\phi)}{(1+\gamma t)^4}.$$

A straightforward calculation now shows that $|R(z_T, z_R)| \le 1$ for all $t \ge 0$ and all ϕ. The result of the lemma then follows from the maximum modulus theorem. □

Lemma 3. No positive angle α exist for which the source-splitting stability function (7.25) is $A(\alpha)$-stable with respect to z_R uniformly for z_T in the stability region of the explicit trapezoidal rule.

Proof. Letting z_R sufficiently small in (7.26) yields

$$R(z_R, z_T) = e^{i\phi} + \tfrac{1}{2} z_R (1 + e^{i\phi}) + O(z_R^2).$$

Substitution of $z_R = a + ib$ with $a \leq 0$ and $e^{i\phi} = \cos\phi + i \sin\phi$ yields

$$|R(z_R, z_T)|^2 = 1 + a(1 + \cos\phi) + b \sin\phi + \tfrac{1}{2}(a^2 + b^2)(1 + \cos\phi) + O(z_R^2).$$

Hence, in first approximation we can write

$$|R(z_R, z_T)|^2 = 1 + a(1 + \cos\phi) + b \sin\phi.$$

Choose $\phi = \pi + \epsilon$ with $\epsilon > 0$ and also arbitrary small. Then, again in first approximation,

$$|R(z_R, z_T)|^2 = 1 - b\epsilon + \tfrac{1}{2} a\epsilon^2 + O(\epsilon^3).$$

We see that $1 - b\epsilon + \frac{1}{2} a\epsilon^2 > 1$ if $b < \frac{1}{2} a\epsilon$. Hence for any sufficiently small real part a, an arbitrarily small imaginary part b exists such that $|R(z_R, z_T)| > 1$, showing that no positive angle α exists defining $A(\alpha)$-stability. □

Numerical experiments with method (7.23) - (7.24) applied to the 3D prototype model using $K_{\max} = 30$ revealed instability. In a similar vein we have studied a source-splitting counterpart of the Strang-splitting method (7.15). For this alternative source-splitting method, using the implicit-explicit transport solver (7.14) instead of the explicit trapezoidal rule, essentially the same restrictive linear stability results do hold as for (7.23) - (7.24). Surprisingly, applied to the 3D problem with the 100 times larger vertical diffusion coefficient, the alternative method based on (7.14) was stable and in fact equally accurate as the approximate factorization ROS2 method (7.19) - (7.20). As yet our findings on source-splitting are therefore inconclusive. Apparently, the precise meaning of the lack of A(α)-stability for z_R must be reconsidered, as well as the use of the implicit-explicit transport solver (7.14) versus the explicit trapezoidal rule.

7.7 Main conclusion

Up to now Strang-type operator splitting seems the method of choice for time stepping in global air quality modeling, in spite of the occurrence of splitting errors. We have found it reliable and it provides flexibility, both for model and code development. Within splitting one of the most time-consuming computations is the stiff chemistry integration. Due to artificial transients introduced at the beginning of split intervals, a highly stable solver is required. This solver should be able to use large time steps in the order of minutes, being far greater than the smallest time constants which are in the order of milliseconds and even smaller. In addition, for convenient code design on vector/parallel and massively parallel computers, this solver should be able

to cope with such large step sizes equally distributed over the whole space grid, or large parts thereof, under widely inhomogeneous spatial and temporal conditions. For three different sets of chemical reactions we have demonstrated that the sparse Rosenbrock solver ROS2 is an excellent candidate. An open question still is how ROS2 will perform under even more difficult real atmospheric and meteorological conditions. Likewise it is of interest to examine the role of the splitting error under such conditions. Finally, ROS2 can of course also prove useful for use in related reactive-flow computations where Strang-splitting is used.

Bibliography

[AB96] I. Ahmad and M. Berzins. An algorithm for ODEs from atmospheric dispersion problems. Technical report, The University of Leeds, 1996. Preprint School of Computer Studies, to appear in Appl. Numer. Math.

[ADRG91] D.J. Allen, A.R. Douglas, R.B. Rood, and P.D. Guthrie. Application of a monotonic upstream-biased transport scheme to three-dimensional constituent transport calculations. *Mon. Wea. Rev.*, 119:2456–2464, 1991.

[BBH89] P.N. Brown, G.D. Byrne, and A.C. Hindmarsh. VODE: A variable coefficient ODE solver. *SIAM J. Sci. Stat. Comput.*, 10:1038–1051, 1989.

[BC89] R.A. Brost and R.B. Chatfield. Transport of radon in a three-dimensional, subhemispheric model. *J. Geophys. Res.*, 94:5095–5119, 1989.

[BHV94] J.G. Blom, W. Hundsdorfer, and J.G. Verwer. Vectorization aspects of a spherical advection scheme on a reduced grid. Report NM-R9418, CWI, Amsterdam, 1994.

[BJ90] Y.J. Balkanski and D.J. Jacob. Transport of continental air to the subantarctic Indian Ocean. *Tellus*, 42B:62–75, 1990.

[Bot92] A. Bott. Monotone flux limitation in the area-preserving flux-form advection algorithm. *Mon. Wea. Rev.*, 120:2595–2602, 1992.

[BR97] J.G. Blom and M.G.M. Roemer. Description of the 3D LOTOS model. Part I: Dynamics. Note MAS-N9701, CWI, Amsterdam, 1997.

[Bre84] Y. Brenier. Averaged multivalued solutions for scalar conservation laws. *SIAM J. Numer. Anal.*, 21:1013–1037, 1984.

[BW76] R.M. Beam and R.F. Warming. An implicit finite-difference algorithm for hyperbolic systems in conservation-law form. *J. Comput. Phys.*, 22:87–110, 1976.

[CJG+96] M. Chin, D.J. Jacob, G.M. Gardner, M.S. Foreman-Fowler, and P.A. Spiro. A global three-dimensional model of tropospheric sulfate. *J. Geophys. Res.*, 101:18667–18690, 1996.

[Cra91] Cray Research, Inc. *Parallel Processing Guide, CF77 Compiling System, Volume 4*, 1991.

[CSP+96] G.R. Carmichael, A. Sandu, F. Potra, V. Damian-Iordache, and M. Damian-Iordache. The current state and the future directions in air quality modeling. In A. Sydow, G.R. Carmichael, and G. Korn, editors, *Modelling and Simulation of Complex Environmental Problems (Proceedings of a Dagstuhl Seminar)*. Springer-Verlag, 1996.

[CZ91] P.J. Crutzen and P.H. Zimmermann. The changing photochemistry of the troposphere. *Tellus*, 43:136–151, 1991.

[DIS95] V. Damian-Iordache and A. Sandu. *KPP - A symbolic preprocessor for chemistry kinetics - User's guide.* University of Iowa, department of Mathematics, 1995.

[DV84] K. Dekker and J.G. Verwer. *Stability of Runge-Kutta methods for stiff nonlinear differential equations.* North-Holland Elsevier Science Publishers, 1984.

[Eas92] R.C. Easter. Two modified versions of Bott's positive-definite numerical advection scheme. *Mon. Wea. Rev.*, 121:297–304, 1992.

[Elb97] H. Elbern. Parallelization and load balancing of a comprehensive atmospheric chemistry transport model, 1997. To appear in Atmos. Environ.

[FKR+96] J.E. Feichter, E. Kjellström, H. Rohde, F.J. Dentener, J. Lelieveld, and G.J. Roelofs. Simulation of the tropospheric sulfur cycle in a global climate model. *Atmos. Environ.*, 30:1693–1708, 1996.

[FRSW92] J.E. Feichter, E. Roeckner, U. Schlese, and M. Windelband. Tracer transport in the Hamburg climate model. In H. van Dop, editor, *Air Pollution Modeling and Its Application VIII.* Plenum Press, New York, 1992.

[GÅU96] R. Gibson, P.K. Ållberg, and S. Uppala. The ECMWF Re-Analysis (ERA) project. *ECMWF Newsletter*, 73:7–17, 1996.

[GM91] J. Graf and N. Moussiopoulos. Intercomparison of two models for the dispersion of chemically reacting pollutants. *Beitr. Phys. Atmos.*, 64:13–25, 1991.

[GvL90] G.H. Golub and C.F. van Loan. *Matrix Computations*. John Hopkins University Press, Baltimore, second edition, 1990.

[GWKD89] M.W. Gery, G.Z. Whitten, J.P. Killus, and M.C. Dodge. A photochemical kinetics mechanism for urban and regional scale computer modeling. *J. Geophys. Res.*, 94:12925–12956, 1989.

[HBCH93] O. Hertel, R. Berkowicz, J. Christensen, and Ø. Hov. Test of two numerical schemes for use in atmospheric transport-chemistry models. *Atmos. Environ.*, 27:2591–2611, 1993.

[Hei95] M. Heimann. The global atmospheric tracer model TM2. Technical Report 10, Deutsches Klimarechnenzentrum (DKRZ), Hamburg, Germany, 1995.

[HHI78] E. Hesstvedt, Ø. Hov, and I.S.A. Isaksen. Quasi-steady state approximation in air pollution modeling: comparison of two numerical schemes for oxidant prediction. *Int. J. Chem. Kinet.*, 10:971–994, 1978.

[HK89] M. Heimann and C.D. Keeling. A three-dimensional transport model for atmospheric CO_2. Part 2: Model description and simulated tracer experiments. In D.H. Peterson, editor, *Aspects of climate variability in the pacific and the Western Americas, Geophysicial monograph 55*, pages 237–275. AGU Washington, 1989.

[HKvLV95] W. Hundsdorfer, B. Koren, M. van Loon, and J.G. Verwer. A positive finite-difference advection scheme. *J. Comput. Phys.*, 117:35–46, 1995.

[HMP90] M. Heimann, P. Monfray, and G. Polian. Modeling the long-range transport of ^{222}Rn to subantarctic and antarctic areas. *Tellus*, 42B:83–99, 1990.

[HS95] W. Hundsdorfer and E.J. Spee. An efficient horizontal advection scheme for modeling of global transport of constituents. *Mon. Wea. Rev.*, 123:3554–3564, 1995.

[HS96] W. Hundsdorfer and E.J. Spee. An atmospheric test model for global advection and chemistry. In E. Kreuzer and O. Mahrenholtz, editors, *Proceedings of the Third International Congress on Industrial and Applied Mathematics (ICIAM/GAMM 95), ZAMM, Issue 4: Applied Sciences, especially Mechanics*, pages 442–443. Akademie Verlag, 1996.

[HT94] W. Hundsdorfer and R.A. Trompert. Method of lines and direct discretization: a comparison for linear advection. *Appl. Num. Math.*, 13:469–490, 1994.

[HV95] W. Hundsdorfer and J.G. Verwer. A note on splitting errors for advection-reaction equations. *Appl. Num. Math.*, 18:191–100, 1995.

[HW96] E. Hairer and G. Wanner. *Solving Ordinary Differential Equations II. Stiff and Differential-Algebraic Problems.* Springer-Verlag, Berlin, second edition, 1996.

[Jc97] D.J. Jacob and 29 co-authors. Evaluation and intercomparison of global atmospheric transport models using ^{222}Rn and other short-lived tracers. *J. Geophys. Res.*, 102:5953–5970, 1997.

[JPWM87] D.J. Jacob, M.J. Prather, S.C. Wofsy, and M.B. McElroy. Atmospheric distribution of ^{85}Kr simulated with a general circulation model. *J. Geophys. Res.*, 92:6614–6626, 1987.

[JSPC97] L.O. Jay, A. Sandu, F.A. Potra, and G.R. Carmichael. Improved quasi-steady-state-approximation methods for atmospheric chemistry integration. *SIAM J. Sci. Comput.*, 18:182–202, 1997.

[JT94] M.Z. Jacobson and R.P. Turco. SMVGEAR: A sparse-matrix, vectorized Gear code for atmospheric models. *Atmos. Environ.*, 28:273–284, 1994.

[Kes95] Chr. Kessler. *Entwicklung eines effizienten Lösungsverfahrens zur modellmäßigen Beschreibung der Ausbreitung und chemischen Umwandlung reaktiver Luftschadstoffe.* PhD thesis, Fakultät für Maschinenbau der Universität Karlsruhe, Verlag Shaker, Aachen, 1995.

[Kor93] B. Koren. A robust upwind discretization method for advection, diffusion and source terms. In C.B. Vreugdenhil and B. Koren, editors, *Numerical Methods for Advection-Diffusion Problems*, Notes on Numerical Fluid Mechanics 45, pages 117–138, Braunschweig, 1993. Vieweg.

[KW93] O. Knoth and R. Wolke. A comparison of fast chemical kinetic solvers in a simple vertical diffusion model. In *Proceedings 20th International Meeting on Air Pollution Modelling and its Applications*, Valencia, Spain, 1993.

[LeV82] R.J. LeVeque. Large time step shock-capturing techniques for scalar conservation laws. *SIAM J. Numer. Anal.*, 19:1091–1109, 1982.

[LeV92] R.J. LeVeque. *Numerical methods for conservation laws.* Lectures in Mathematics. ETH Zürich, Birkhäuser Verlag, Basel, second edition, 1992.

[Lou79] J.-F. Louis. A parametric model of vertical eddy fluxes in the atmosphere. *Bound.-Layer Meteor.*, 17:187–202, 1979.

[Lou86] J.-F. Louis. Research manual 2; ECMWF forecast model; adiabatic part. Technical report, ECMWF Research Department, 1986.

[LY90] R.J. LeVeque and H.C. Yee. A study of numerical methods for hyperbolic conservation laws with source terms. *J. Comput. Phys.*, 86:187–210, 1990.

[Mat95] J. Matthijssen. Laboratoire d'Aerologie OMP, Toulouse, France, 1995. Private Communication.

[MB95] J.-F. Müller and G. Brasseur. IMAGES: A three-dimensional chemical transport model of the global troposphere. *J. Geophys. Res.*, 100:16445–16490, 1995.

[MGS82] G.J. McRae, W.R. Goodin, and J.H. Seinfeld. Numerical solution of the atmospheric diffusion equation for chemically reacting flows. *J. Comput. Phys.*, 45:1–41, 1982.

[PBC+95] L.K. Peters, C.M. Berkowitz, G.R. Carmichael, R.C. Easter, G. Fairweather, S.J. Ghan, J.M. Hales, L.R. Leung, W.R. Pennell, F.A. Potra, R.D. Saylor, and T.T. Tsang. The current state and future directions of Eulerian models in simulating the tropospheric chemistry and transport of trace species: a review. *Atmos. Environ.*, 29:189–222, 1995.

[Phi57] N.A. Phillips. A coordinate system having some special advantages for numerical forecasting. *J. Meteor.*, 14:184–185, 1957.

[PMW+87] M.J. Prather, M.B. McElroy, S.C. Wofsy, G. L. Russell, and D. Rind. Chemistry of the global troposphere: Fluorocarbons as tracers of air motion. *J. Geophys. Res.*, 92:6579–6613, 1987.

[Pra86] M.J. Prather. Numerical advection by conservation of second-order moments. *J. Geophys. Res.*, 91:6671–6681, 1986.

[Pro97] Proceedings Saint-Venant Symposium, Paris. *Reducing mechanism in chemical kinetics for the simulation of reactive transport: an application to air pollution modeling*, 1997.

[Ras94] P.J. Rasch. Conservative shape-preserving two-dimensional transport on a spherical reduced grid. *Mon. Wea. Rev.*, 122:1337–1350, 1994.

[RL81] G.L. Russell and J.A. Lerner. A new finite-differencing scheme for the tracer transport equation. *J. Appl. Meteor.*, 20:1483–1498, 1981.

[RLSK97] G.J. Roelofs, J. Lelieveld, H.G.J. Smit, and D. Kley. Ozone production and transports in the tropical Atlantic region during the biomass burning season. *J. Geophys. Res.*, 102:10637–10651, 1997.

[Roo87] R.B. Rood. Numerical advection algorithms and their role in atmospheric transport and chemistry models. *Rev. Geophys.*, 25:71–100, 1987.

[SdZV+96] E.J. Spee, P.M. de Zeeuw, J.G. Verwer, J.G. Blom, and W. Hundsdorfer. Vectorization and parallelization of a numerical scheme for 3D global atmospheric transport-chemistry problems. Report NM-R9620, CWI, Amsterdam, 1996.

[Sim93] D. Simpson. Photochemical model calculations over Europe for two extended summer periods: 1985 and 1989. Model results and comparisons with observations. *Atmos. Environ.*, 27A:921–943, 1993.

[Smo82] P.K. Smolarkiewicz. The multi-dimensional Crowley advection scheme. *Mon. Wea. Rev.*, 110:1968–1983, 1982.

[SPCD96] 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.

[Spe95] E.J. Spee. Coupling advection and chemical kinetics in a global atmospheric test model. In H. Power, N. Moussiopoulos, and C.A. Brebbia, editors, *Air Pollution III, Volume 1: Air Pollution, Theory and Simulation*, pages 319–326. Computational Mechanics Publications, Southampton, Boston, 1995.

[SPvDH97] E.J. Spee, A.C. Petersen, H. van Dop, and W. Hundsdorfer. Sensitivity of atmospheric transport model performance to numerical advection schemes and resolution. Report MAS-R9710, CWI, Amsterdam, 1997. Submitted to J. Geophys. Res.

[SR91] P.K. Smolarkiewicz and P.J. Rasch. Monotone advection on the sphere: an Eulerian versus semi-Lagrangian approach. *J. Atmos. Sci.*, 48:793–810, 1991.

[Str62] G. Strang. Trigonometric polynomials and difference methods of maximal accuracy. *J. Math. and Phys.*, 41:147–154, 1962.

[Str68] G. Strang. On the construction and comparison of difference schemes. *SIAM J. Numer. Anal.*, 5:506–517, 1968.

[Sun96] P. Sun. A pseudo-non-time-splitting method in air quality modeling. *J. Comput. Phys.*, 127:152–157, 1996.

[SVB+97] A. Sandu, J.G. Verwer, J.G. Blom, E.J. Spee, and G.R. Carmichael. Benchmarking stiff ODE solvers for atmospheric chemistry problems II: Rosenbrock solvers. *Atmos. Environ.*, 31:3459–3472, 1997.

[SVdZ+97a] E.J. Spee, J.G. Verwer, P.M. de Zeeuw, J.G. Blom, and W. Hundsdorfer. A numerical study for global atmospheric transport-chemistry problems, 1997. Revision of MAS-R9702.

[SVdZ+97b] E.J. Spee, J.G. Verwer, P.M. de Zeeuw, J.G. Blom, and W. Hundsdorfer. A numerical study for global atmospheric transport-chemistry problems. Report MAS-R9702, CWI, Amsterdam, 1997.

[SVvL+97] A. Sandu, J.G. Verwer,, M. van Loon, G.R. Carmichael, F.A. Potra, D. Dabdub, and J.H. Seinfeld. Benchmarking stiff ODE solvers for atmospheric chemistry problems I: implicit versus explicit. *Atmos. Environ.*, 31:3151–3166, 1997.

[Swe84] P.K. Sweby. High resolution schemes using flux-limiters for hyperbolic conservation laws. *SIAM J. Numer. Anal.*, 21:995–1011, 1984.

[The94] H. The. RIVM, Bilthoven, The Netherlands, 1994. Private Communication.

[Tie89] M. Tiedtke. A comprehensive mass flux scheme for cumulus parametrization in large-scale models. *Mon. Wea. Rev.*, 117:1779–1800, 1989.

[VB96] J.G. Verwer and J.G. Blom. On the coupled solution of diffusion and chemistry in air pollution models. In E. Kreuzer and O. Mahrenholtz, editors, *Proceedings of the Third International Congress on Industrial and Applied Mathematics (ICIAM/GAMM 95), ZAMM, Issue 4: Applied Sciences, especially Mechanics*, pages 454–457. Akademie Verlag, 1996.

[VBH96] J.G. Verwer, J.G. Blom, and W. Hundsdorfer. An implicit-explicit approach for atmospheric transport-chemistry problems. *Appl. Num. Math.*, 20:191–209, 1996.

[VBvLS96] J.G. Verwer, J.G. Blom, M. van Loon, and E.J. Spee. A comparison of stiff ODE solvers for atmospheric chemistry problems. *Atmos. Environ.*, 30:49–58, 1996.

[vdHSK96] P.J. van der Houwen, B.P. Sommeijer, and J. Kok. The iterative solution of fully implicit discretizations of three-dimensional transport models. Report NM-R9621, CWI, Amsterdam, 1996.

[Ver77] J.G. Verwer. S-stability properties for generalized Runge-Kutta methods. *Numer. Math.*, 27:359–370, 1977.

[Ver94] J.G. Verwer. Gauss-Seidel iteration for stiff ODEs from chemical kinetics. *SIAM J. Sci. Comput.*, 15:1243–1250, 1994.

[VHK94] G.J.M. Velders, L.C. Heijboer, and H. Kelder. The simulation of the transport of aircraft emissions by a three-dimensional global model. *Ann. Geophys.*, 12:385–393, 1994.

[VK93] C.B. Vreugdenhil and B. Koren, editors. *Numerical Methods for Advection-Diffusion Problems*, volume 45 of *Notes on Numerical Fluid Mechanics*. Vieweg, Braunschweig, 1993.

[vL77] B. van Leer. Towards the ultimate conservative difference scheme IV, a new approach to numerical convection. *J. Comput. Phys.*, 23:276–299, 1977.

[vL96] M. van Loon. *Numerical methods in smog prediction.* PhD thesis, Universiteit van Amsterdam, 1996.

[VS95] J.G. Verwer and D. Simpson. Explicit methods for stiff ODEs from atmospheric chemistry. *Appl. Num. Math*, 18:413–430, 1995.

[VSBH97] J.G. Verwer, E.J. Spee, J.G. Blom, and W. Hundsdorfer. A second-order Rosenbrock method applied to photochemical dispersion problems. Report MAS-R9717, CWI, Amsterdam, 1997. Submitted to SIAM J. Sci. Comput.

[War88] P. Warneck. *Chemistry of the natural atmosphere.* Academic Press, Inc., San Diego, 1988.

[WB79] R.F. Warming and R.M. Beam. An extension of A-stability to alternating direction methods. *BIT*, 19:395–417, 1979.

[Wil92] D.L. Williamson. Review of numerical approaches for modeling global transport. In H. van Dop and G. Kallos, editors, *Air Pollution Modeling and Its Application IX*, pages 377–394. NATO Challenges of Modern Society 17, Plenum Press, New York, 1992.

[WR89] D.L. Williamson and P.J. Rasch. Two-dimensional semi-Lagrangian transport with shape-preserving interpolation.
Mon. Wea. Rev., 117:102–129, 1989.

[Zal87] S.T. Zalesak. A preliminary comparison of modern shock-capturing schemes: linear advection. In R. Vichnevetsky and R.S. Stapelman, editors, *Advances in Computer Methods for Partial Differential Equations IV*, pages 15–22. IMACS, 1987.

[Zim84] P.H. Zimmermann. *Ein dreidimensionales numerisches Transportmodell für atmosphärische Spurenstoffe.* PhD thesis, Fachbereich Physik der Johannes Gutenberg - Universität, Mainz, 1984.

[Zim87] P.H. Zimmermann. MOGUNTIA: A handy global tracer model. In H. van Dop, editor, *Air Pollution modeling and its application VI*, pages 593–608. Plenum, New York, 1987.

[Zla95] Z. Zlatev. *Computer treatment of large air pollution models.* Kluwer Academic Publishers, 1995.

Samenvatting

Numerieke methoden in globale transport-chemie modellen

Het onderwerp van dit proefschrift betreft de numerieke wiskunde in globale transport-chemie modellen. Met deze modellen wordt onderzoek verricht naar de chemische samenstelling van de atmosfeer, in het bijzonder in relatie tot luchtverontreining op globale schaal. Met globale modellen wordt derhalve onderzoek verricht naar globale milieuproblemen, zoals het ozongat boven de Zuidpool en het antropogeen (door menselijke activiteiten) versterkte broeikaseffect, dat mogelijk resulteert in globale opwarming. Een gerelateerd onderzoeksgebied is smogvoorspelling. Met dit verschil dat smogvorming voornamelijk plaats vindt op een veel kleinere schaal, bijvoorbeeld in de omgeving van stedelijke gebieden.

Het in dit proefschrift beschreven onderzoek is uitgevoerd in het kader van het CIRK-project. Het CIRK-project is ontstaan uit contacten tussen CWI, IMAU, RIVM en KNMI met als doel de verbetering van globale 3D modellering van de troposfeer, waarin chemie, emissies, deposities en troposfeer-stratosfeer uitwisseling zijn opgenomen. Dit proefschrift beschrijft numerieke algoritmen die in deze modellen voorkomen. Op het IMAU wordt parallel aan dit onderzoek gewerkt aan vertikaal transport en processen op sub-grid schaal. Voor het onderzoek in dit proefschrift is het contact met wetenschappers van IMAU, RIVM en KNMI zeer belangrijk geweest.

We beschouwen transport-chemie modellen die zijn gebaseerd op balansvergelijkingen. De chemische stoffen in deze modellen hebben zeer verschillende tijdschalen: van minder dan een seconde voor radicalen tot jaren voor bijvoorbeeld methaan. De vergelijkingen worden opgelost op een Euleriaans lengtegraad-breedtegraad rooster. Het product van het aantal roosterpunten en het aantal chemische componenten ligt tussen de honderdduizend en één miljard. De combinatie van het grote aantal onbekenden en de verschillende tijdschalen maakt het oplossen van deze vergelijkingen zeer rekenintensief. Er worden dan ook hele hoge eisen gesteld aan zowel de computers als aan de numerieke algoritmen.

Vanwege de grootte van modelfouten, bijvoorbeeld in de emissies en de chemische reactiesnelheden, volstaat in deze modellen een lage nauwkeurigheid. Het model kan daarom worden opgelost met operator splitting. Dit heeft tot gevolg dat voor processen als advectie en chemie verschillende numerieke technieken gebruikt kunnen worden. In dit proefschrift ligt de nadruk op numerieke technieken voor advectie en chemie en alternatieven voor de veelgebruikte operator splitting.

Bij de ontwikkeling van advectieschema's hebben we ons gericht op massa behoudende schema's die geen nieuwe minima of maxima introduceren, zogenaamde monotone schema's. Massabehoud wordt afgedwongen door advectie te beschrijven met in- en uitgaande fluxen. Monotoniciteit wordt verkregen door het gebruik van

flux-limiters.

Op de bol geeft het gebruikte rooster aanleiding tot een singulariteit aan de beide polen. Er zijn twee mogelijkheden onderzocht om de bezwaren van deze singulariteit te onderdrukken: het zogenaamde gereduceerde rooster waar aan de polen minder cellen worden gebruikt, en een onvoorwaardelijk stabiel schema. Goede resultaten werden verkregen met het schema Split-rg. Dit schema maakt gebruik van dimensie-splitting op een gereduceerd rooster. In vergelijking tot bestaande schema's is Split-rg nauwkeurig, geheugen-efficiënt en gebruikt het weinig cpu-tijd. Het schema Mol-rg geeft ook bevredigende resultaten. Mol-rg is een Methode der Lijnen schema op een gereduceerd rooster, dat gebruik maakt van derde orde upwind fluxen met flux-limiters en de expliciete trapeziumregel voor tijdsintegratie.

Het meest rekenintensieve deel in deze modellen is het oplossen van de chemie. Een reductie van de cpu-tijd kan worden bereikt door grote tijdstappen te nemen. Dit is echter alleen mogelijk als de chemie solver goede stabiliteitseigenschappen bezit. Het tweede orde Rosenbrock schema ROS2 lijkt heel geschikt als chemie solver. Met ROS2 kunnen we een complexe chemie integreren met een tijdstap van 20 minuten, zonder gebruik te maken van enige voorkennis van de gebruikte chemie.

De veel toegepaste techniek operator splitting heeft een aantal grote voordelen. De verschillende processen kunnen met verschillende technieken worden opgelost. De advectie kan expliciet en de chemie impliciet worden opgelost. Tevens kan de chemie per roosterpunt in tijdstap variëren. Verder is de computercode modulair van opzet en geheugen-efficiënt. Helaas ontstaat bij operator splitting een splitfout. We hebben geconcludeerd dat in globale modellen de splitfout binnen de gewenste nauwkeurigheid blijft. Hier is echter meer onderzoek nodig. Als alternatief voor operator splitting hebben we twee methoden bestudeerd om chemie en vertikale diffusie gekoppeld op te lossen: Twostep, een tweede orde BDF methode (Backward differentiation formula) met een speciale Gauss-Seidel iteratie en de gefactoriseerde ROS2. ROS2 met factorisatie is een goed alternatief voor operator splitting, maar met Twostep kan helaas natte chemie niet worden opgelost.

De ontwikkelde technieken zijn uitgebreid getest. Daarvoor zijn standaard testen gebruikt zoals de Molenkamp test op een bol en state-of-the-art chemie boxmodellen. Ook zijn deze boxmodellen gecombineerd met de Molenkamp test om zo voor een complete advectie-chemie-diffusie stap de nauwkeurigheid tegen de cpu-tijd af te zetten.

Dit proefschrift behandelt de numerieke wiskunde in globale transport-chemie modellen.
Deze modellen worden gebruikt om onderzoek te verrichten naar onder andere het gat in de ozonlaag en het broeikaseffect.
In dit proefschrift worden twee onderdelen uit het model uitvoerig behandeld: advectie op een bol en stijve differentiaal-vergelijkingen afkomstig van atmosferische chemie.
Advectie op een bol wordt opgelost met een eindige-volume-discretisatie op een gereduceerd rooster. De chemie wordt met Rosenbrock solvers opgelost.
Het onderzoek is verricht op het CWI te Amsterdam.