



A mass-conserving and multi-tracer efficient transport scheme in the online integrated Enviro-HIRLAM model

B. Sørensen¹, E. Kaas¹, and U. S. Korsholm²

¹Niels Bohr Institute, University of Copenhagen, Copenhagen, Denmark

²Danish Meteorological Institute, Copenhagen, Denmark

Correspondence to: B. Sørensen (sorensen@gfy.ku.dk)

Received: 23 October 2012 – Published in Geosci. Model Dev. Discuss.: 20 November 2012

Revised: 7 June 2013 – Accepted: 14 June 2013 – Published: 22 July 2013

Abstract. In this paper a new advection scheme for the online coupled chemical–weather prediction model Enviro-HIRLAM is presented. The new scheme is based on the locally mass-conserving semi-Lagrangian method (LMCSL), where the original two-dimensional scheme has been extended to a fully three-dimensional version. This means that the three-dimensional semi-implicit semi-Lagrangian scheme which is currently used in Enviro-HIRLAM is largely unchanged. The HIRLAM model is a computationally efficient hydrostatic operational short-term numerical weather prediction model, which is used as the base for the online integrated Enviro-HIRLAM. The new scheme is shown to be efficient, mass conserving, and shape preserving, while only requiring minor alterations to the original code. It still retains the stability at long time steps, which the semi-Lagrangian schemes are known for, while handling the emissions of chemical species accurately. Several mass-conserving filters have been tested to assess the optimal balance of accuracy vs. efficiency.

1 Introduction

In numerical weather prediction (NWP) and atmospheric chemical transport (ACT) models, the numerical methods used for advection are chosen based upon the specific use of the model, e.g. short-term weather prediction, long-term climate prediction, air pollution modelling, or emergency preparedness models. A series of desirable properties for transport schemes have been suggested (Rasch and Williamson, 1990; Machehauer et al., 2007), e.g. accuracy, conservation, and shape preservation. In any specific application some of

the desirable properties are more important than others, and the less important ones are often not addressed for the sake of efficiency. Combining models, as is done in integrated chemical weather forecasting (CWF) (Lawrence et al., 2005), the original less important properties might not be negligible any more or the transport scheme used becomes insufficient for the new problem, and the combined numerical schemes should be revised (Baklanov, 2008, 2010).

In this paper the Enviro-HIRLAM model, an online integrated chemistry extension to the High Resolution Limited Area Model (HIRLAM), has inherited the numerical properties of the HIRLAM model (Korsholm et al., 2008; Korsholm, 2009). The HIRLAM model is a well proven and numerically efficient hydrostatic NWP model; however, some of the properties originally considered negligible have an impact on ACT modelling, as shown in this paper. Other models, which are under the same kind of development, are facing similar challenges. For example, the Chemical – Integrated Forecast System (C-IFS) at the European Centre for Medium Range Weather Forecast (ECMWF) has implemented additional measures to ensure global mass conservation for the chemical species (Flemming et al., 2009, 2012).

The properties of mass conservation, shape preservation, and multi-tracer efficiency are important for online integrated NWP/ACT models (Thuburn, 2008). Mass conservation is of importance since spurious numerical sources and sinks of chemical species reduce the forecast credibility. Shape preservation is also important since non-linear chemistry introduces spurious reactions initiated solely by the deficiencies of the numerical scheme. Here it is important to note that shape preservation should be fulfilled in mixing ratio as it is the relative concentrations of the individual chemical

species that determine if any chemical reactions are triggered (Thuburn and McIntyre, 1997; Lauritzen et al., 2006). The term multi-tracer efficiency is related to the amount of additional cost for each additional advected tracer¹. When introducing tens or hundreds of new tracers, it is crucial that the scheme can transport the tracers with as little computational overhead as possible.

The Enviro-HIRLAM model was originally constructed with a separate advection method, the Bott scheme (Bott, 1989), for the chemical species only (Chenevez et al., 2004). This was done to ensure mass conservation. The Bott scheme is, however, an Eulerian method which suffers from the Courant–Friedrich–Lewy (CFL) criterion, thus needing multiple sub-steps for each semi-Lagrangian (SL) time step. This, combined with very little multi-tracer efficiency, makes the Bott transport scheme unsuitable for operational models.

The aims of this study are to implement and validate the new transport scheme and to assess the importance of the desirable properties, which many NWP/ACT models neglect to some extent. Another additional property is also addressed, namely internal consistency, which is often not considered a priority, but it is nevertheless a property which should be met if possible. This implies that all density variables are transported using the same method (Jockel et al., 2001), including the prognostic variables of the meteorological model (e.g. specific humidity, cloud water, and cloud ice).

The new scheme is essentially an extension of the original semi-Lagrangian scheme, comprised of a modification of the advected fields ensuring mass conservation and the addition of a locally mass-conserving and shape-preserving filter.

The Enviro-HIRLAM model is briefly described in the next section. The LMCSL scheme is introduced in Sect. 3, then in Sect. 4 different filter methods are considered, and in Sect. 5 the results are presented and discussed. The final section contains some concluding remarks.

2 Enviro-HIRLAM

The work for the Enviro(nment)-HIRLAM model in its present form was initiated in 2004 at the Danish Meteorological Institute (DMI) with the implementation of the Bott advection scheme for passive tracer transport in the DMI-HIRLAM model (Chenevez et al., 2004; Korsholm et al., 2008). In recent years, development of new and improved transport schemes for the model has been carried out. In 2008 a mass-conserving semi-Lagrangian scheme, the so-called cell-integrated semi-Lagrangian scheme (CISL), was implemented (Lauritzen et al., 2008). The CISL scheme was, however, never fully adopted as the default dynamical core in the model. Later, another semi-Lagrangian mass-conserving scheme (Sørensen et al., 2012), developed from the CISL-HIRLAM model, was implemented. This scheme

¹Tracer is here used as a broad term meaning any advected prognostic density variable.

was using the two-dimensional locally mass-conserving semi-Lagrangian (LMCSL) method (Kaas, 2008) in order to acquire mass conservation. Both the CISL and the LMCSL-2D scheme require quite large changes in the underlying model code.

The Enviro-HIRLAM model is now a part of the HIRLAM framework as the chemical branch for HIRLAM (Baklanov et al., 2011). The model utilizes a fully online integrated chemical scheme, and earlier versions included two-way feedbacks between pollutants and meteorological processes (Korsholm, 2009; Baklanov, 2010). The chemical scheme currently implemented is a condensed version of the CBM-Z scheme (Zaveri and Peters, 1999).

3 Locally mass-conserving semi-Lagrangian scheme

The LMCSL method has previously been implemented into the Enviro-HIRLAM model (as mentioned above) in a quasi-Lagrangian (cascade-like) manner in which the advection is split into a horizontal part and a vertical part. The quasi-Lagrangian version required substantial changes to the original model code, which makes further development from the HIRLAM system quite difficult and time consuming. Here a fully three-dimensional implementation (LMCSL-3D) is used instead. This reduces the code complexity considerably as the changes are limited to a few additional routines, which are compatible with the official HIRLAM code. The LMCSL-3D scheme uses the same approach as the two-dimensional LMCSL scheme to achieve mass conservation. The basics are explained below; see Kaas (2008) for further details.

In a semi-Lagrangian (SL) model the Lagrangian form of the continuity equation is used to describe the evolution of the volume density. If ψ is the volume density and \mathbf{v} is the velocity, then the continuity equation for ψ becomes

$$\frac{d\psi}{dt} = -\psi \nabla \cdot \mathbf{v}. \quad (1)$$

A two-time-level semi-Lagrangian discretization for integrating (Eq. 1) forward in time can be written as

$$\psi_k^{n+1}{}_{\text{SL-exp}} = \{\psi - 0.5\Delta t(\psi \nabla \cdot \mathbf{v})\}_{*k}^n - 0.5\Delta t(\psi \nabla \cdot \tilde{\mathbf{v}})_k^{n+1}, \quad (2)$$

with k being an Eulerian (arrival) grid point index, Δt is the length of the time step, $(\cdot)_*$ denotes a value that has been interpolated to an upstream departure point, and $(\tilde{\cdot})$ denotes a value that has been extrapolated from time levels $n-1$ and n . The traditional explicit SL scheme (Eq. 2) can be rewritten as

$$\psi_k^{n+1}{}_{\text{SL-exp}} = \sum_{l=1}^K w_{k,l} \left[\psi_l^n - \frac{\Delta t(\psi \nabla \cdot \mathbf{v})_l^n}{2} \right] - \frac{\Delta t(\psi \nabla \cdot \tilde{\mathbf{v}})_k^{n+1}}{2}, \quad (3)$$

where K is the total number of Eulerian grid points, and $w_{k,l}$ the interpolation weight from an upstream Eulerian grid point l (at time level n) to the Lagrangian departure point $(\cdot)_{*k}$. Thus, the $w_{k,l}$ values define the SL remapping required for each Eulerian arrival point. These interpolation weights are not dependent on the actual field values, and are only calculated once and then reused.

Modified interpolation weights

As shown in Kaas (2008) mass conservation can be achieved by modifying the interpolation weights, $w_{k,l}$, and skipping the divergence terms in Eq. (3). The modified weights, $\hat{w}_{k,l}$, are defined as

$$\hat{w}_{k,l} = \frac{V_l}{V_k} \frac{w_{k,l}}{\sum_{m=1}^K w_{m,l}}, \quad (4)$$

which gives

$$\bar{\psi}_k^{n+1} = \sum_{l=1}^K \hat{w}_{k,l} \cdot \bar{\psi}_l^n, \quad (5)$$

with $w_{k,l}$ being the traditional SL weights. V_l and V_k are the volume represented by the l -th and k -th Eulerian grid cells, respectively. Note that the LMCSL-3D scheme forecasts the grid cell averages, here indicated with a bar, and not the grid point values as in the traditional SL scheme. It can easily be shown that, with this approach, the modified weights ensure that the total mass given off by any Eulerian grid cell to all surrounding departure points is equal to the actual mass represented by that particular grid cell. To calculate the modified weights, one loops through all grid cells adding the SL weights, thus obtaining the denominator in Eq. (4). When all modified weights have been calculated, they can be reused for any prognostic density variable in the same way as the SL weights. This makes the method equally multi-tracer efficient since each additional variable does not require any additional computation of interpolation weights. To compute the $w_{k,l}$ weights the traditional HIRLAM interpolation scheme is used. The integration of the LMCSL-3D method into the model is described below:

1. After trajectories are calculated and used to determine the traditional SL interpolation weights, $w_{k,l}$, these weights are passed into a routine that computes the total weight on each grid cell and calculates the normalization factor, $\sum_{m=1}^K w_{m,l}$. This factor is generally close to 1 but can vary between 0.9 and 1.1 near mountains when high Courant numbers are used.
2. The departure cell volumes, V_l , are calculated and divided by the normalization factor, giving us

$$\chi_l = \frac{V_l}{\sum_{m=1}^K w_{m,l}}. \quad (6)$$

3. All transported variables are multiplied by χ_l , and then advected using the traditional advection routine where the multiplication by the SL weights, $w_{k,l}$, are performed.
4. The remainder of the traditional semi-implicit SL scheme is now performed to give us the updated pressure field at time $n + 1$.
5. Shape preserving maximum and minimum values, ψ_k^+ and ψ_k^- , are calculated using the updated pressure field (see details in Sect. 4).
6. The tracer fields are then filtered using the maximum and minimum values. The filter can be either global or local (see details in Sect. 4).
7. Finally, the advected and filtered tracer fields are divided with the arrival volumes, thus concluding the LMCSL advection step, giving us $\bar{\psi}_k^{n+1}$.

The traditional SL scheme includes several interpolation algorithms depending on the desired accuracy, e.g. trilinear, triquadratic, and tricubic interpolation. By using the same algorithms the method minimizes the code complexity as well as enhances the consistency by using the same interpolation algorithm for non-conserved variables, e.g. temperature and velocity fields. The LMCSL-3D scheme thus has the same locality properties as the original SL scheme. With tricubic interpolation, 64 grid cell values are used to determine the value at any departure point. Note that unlike the quasi-Lagrangian LMCSL scheme, the new LMCSL-3D scheme is not advecting the mass field and the mass–wind inconsistency is still present.

3.1 Shape preservation

When dealing with an additional large number of chemical species, some properties are not just desirable but even become essential. In Sect. 3, the issue of mass conservation was considered, but since the scheme uses higher-order interpolations (quadratic and cubic), it will generate both undershoots (containing both positive and negative values) and overshoots. Here under- and overshoots are assumed strictly non-realistic, in the sense that the mixing ratio in a given Lagrangian departure volume cannot be larger than or smaller than the mixing ratio of the Eulerian cells in which it is located.

When using the term shape preservation it is important to remember that it is the mixing ratio that must be preserved. Since chemical reactions are triggered by changes in relative concentrations of the individual constituents, a shape-preserving method will reduce unphysical reactions. Unless one is using fully Lagrangian schemes it is in general not possible to completely eliminate unphysical chemical reactions (Lauritzen et al., 2006).

There are several methods one may use to handle these issues; the one we have chosen is to use a posteriori filtering which can have several different implementations. In the quasi-Lagrangian LMCSL scheme a locally anti-diffusive filter as well as simpler non-anti-diffusive filter was used. The anti-diffusive filter (Kaas and Nielsen, 2008) was only developed for use in two dimensions and is difficult to utilize efficiently on a parallel computing system. It is therefore not used here. Instead a three-dimensional version of the simpler filter, called the iterative locally mass-conserving (ILMC) filter, is implemented and compared to other filters, which have been and still are used in other models.

4 Filters

Negative values are undesirable since they are unphysical and must be removed prior to any chemistry calculations. One can use positive definite methods to handle this, e.g. truncation or redistribution of mass. However, unless positive under- and overshoots are treated, so that shape preservation is fulfilled, the oscillations can have a considerable impact on chemical reactions (Thuburn and McIntyre, 1997).

When using shape-preserving filters the stencil used for advection must be taken into account. If an advection scheme is global in nature, e.g. spectral models, it is generally not possible to make local redistributions which are both mass conserving and shape preserving, since the cause for the under- or overshoot must be in the domain of redistribution; that is, for more local schemes using e.g. tricubic interpolations, the redistribution stencil should be of at least the same size.

In Sects. 4.1 to 4.3 below, several filters which have been implemented are described. The filters are the default iterative locally mass-conserving filter, the local departure-dependent filter, and two global filters – one shape preserving and one positive definite.

4.1 ILMC filter

The iterative locally mass-conserving (ILMC) filter is an a posteriori filter which detects violations of shape preservation and consequently redistributes the mass to the surrounding cells, where the subscripts k and l represent the violated cell and the surrounding cells, respectively. The steps performed in the filter are as follows:

1. Compute shape-preserving maximum and minimum values (entire field),

$$m_k^+ = m_{\text{tot}} \cdot \max(\zeta_{l1}, \dots, \zeta_{l8}), \quad (7)$$

$$m_k^- = m_{\text{tot}} \cdot \min(\zeta_{l1}, \dots, \zeta_{l8}), \quad (8)$$

where $l_1 - l_8$ are the eight Eulerian grid points surrounding to the Lagrangian departure point $(-)*_k$ in

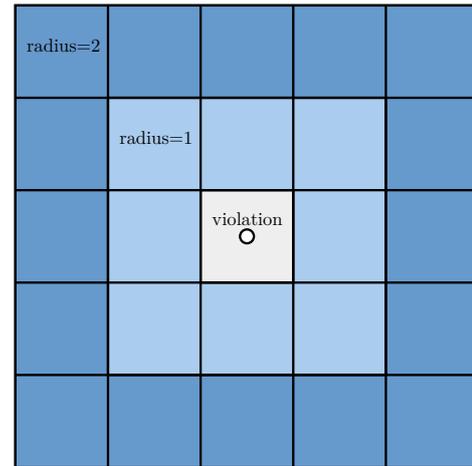


Fig. 1. Illustration of the ILMC filter in two dimensions. The ring in the centre indicates the cell containing a shape-preserving violation. The coloured squares are the cells included in the first (light blue) and second (dark blue) iteration, respectively. In three dimensions the coloured rings become shells consisting of cells in the layers above and below as well.

three-dimensional space, m is the mass, and ζ the mixing ratio.

2. Detect undershoots and overshoots.
3. At violations, compute the mass, Δm_k , available in the surrounding cells.

$$\delta m_l^+ = \max(m_l^+ - m_l, 0), \quad (9)$$

$$\delta m_l^- = \max(m_l - m_l^-, 0), \quad (10)$$

$$\Delta m_k^+ = \sum_{l=1}^{L_r} \delta m_l^+, \quad (11)$$

$$\Delta m_k^- = \sum_{l=1}^{L_r} \delta m_l^-, \quad (12)$$

with L_r being the number of surrounding grid cells at radius r .

4. Add or subtract the amount of available mass needed from the surrounding cells to acquire mass conservation.
5. Repeat from step 3 with the cells in the next *shell*, if available mass was insufficient (see Fig. 1).

Generally one or two iterations are sufficient to redistribute the mass, but up to three iterations are allowed to ensure mass conservation. The ILMC filter is now the default filter used in the mass-conserving Enviro-HIRLAM, and it is used if no other filter is mentioned explicitly in the text.

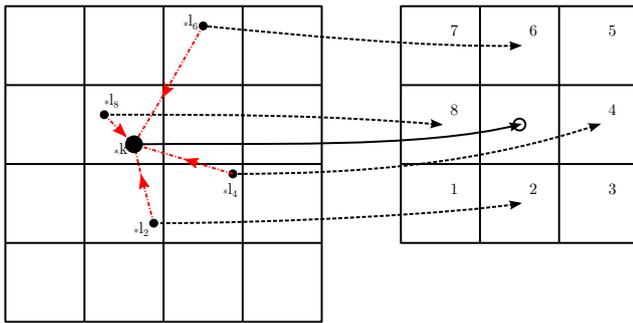


Fig. 2. Departure dependent filter in two dimensions: by redistributing mass to each neighbouring cell one by one, starting with the closest, the mass is distributed as locally as possible. For simplicity the trajectories of the four corner cells [1, 3, 5, 7] are not shown

4.2 Departure dependent filter

The departure-dependent filter (DEPDEP) is implemented as a benchmark for the other filters. The DEPDEP filter is, like ILMC, an iterative filter. All departure points, $*k$, arrive in an Eulerian cell surrounded by eight neighbouring cells² (see Fig. 2). The basic assumption is then that as the filter attempts to redistribute mass to these surrounding cells, not all neighbours are equal, since their departure points were generally not equally distributed. By redistributing the mass based on the departure point distances we essentially get something that increases locality by mimicking a small redistribution in departure point space.

In Fig. 2, the distance between departure point $*k$ and the departure point $*l_8$ is shorter than the distances from $*k$ to $*l_2$, $*l_4$, and $*l_6$. The filter routine is as follows:

- 1–2. Same as for the ILMC filter.
3. At violations, compute distance from the departure point of the undershoot, $*k$, to the departure points of the surrounding cells $*l_{1-8}$.

$$d_{k,l} = \sqrt{(x_{*k}^n - x_{*l}^n)^2 + (y_{*k}^n - y_{*l}^n)^2 + (\tilde{p}_{*k}^n - \tilde{p}_{*l}^n)^2}, \quad (13)$$

where $d_{k,l}$ is the distance, x and y are the horizontal coordinates, and \tilde{p} is the vertical pressure coordinate, where the $(\tilde{\cdot})$ indicates that the pressure coordinate has been normalized for the comparison with the horizontal distances.

4. Sort the surrounding points, shortest distance first.
5. Compute mass available in nearest cell l_1 .

$$\delta m_l^+ = \max(m_l^+ - m_l, 0) \quad (14)$$

$$\delta m_l^- = \max(m_l - m_l^-, 0) \quad (15)$$

²Described here in two dimensions for simplicity

6. Add or subtract the amount of available mass needed from l_1 to acquire mass conservation.
7. Repeat from (5) with next nearest cell, l_{2-8} , if available mass was insufficient.
8. If the available mass is still insufficient, repeat from step 3 with the cells in the next shell.

Since all departure points are known, the distances are trivial to compute. The majority of the computation time is spent sorting the cells according to distance which, when extending it to three dimensions and increasing the radius of redistribution, becomes a rather intensive task (but easily implementable). The DEPDEP filter is used as reference filter in Sect. 5 because it is highly local, while distributing mass in a physically meaningful manner, and shape preserving at the same time. The reason it is not used as the default filter is because it is too computationally expensive in an operational setup.

4.3 Global filters

Globally mass-conserving filters such as the Bartnicki filter (Bartnicki, 1990) have often been used in ACT modelling. The traditional Bartnicki filter is an iterative positive definite global filter, redistributing total negative mass equally among all cells with positive mass. It is repeated if new negative values are generated. Positive definite filters are not well suited for ACT modelling if the methods used generate oscillations (such as most higher order methods), since most oscillations are non-negative, and thus not treated. The filter can, however, quite easily be transformed into a non-iterative shape-preserving filter by not allowing new over- and undershoots to be generated. The Bartnicki-type filters are all extremely fast and trivial to implement both on shared and distributed memory systems.

The reason that global filters are tested here is that they have been used extensively in models, but their global nature might become an issue when dealing with non-linear chemistry, where they can introduce unphysical chemical reactions and possibly mask the effect of small emission changes. Two different global filters are implemented: the shape preserving Global-M and the positive definite Global-PD.

So-called global mass fixers are similar to the above-mentioned filters, which assume that a mass-conserving transport scheme has been used, where the mass fixers are used to compensate for a non-mass-conserving transport scheme. The C-IFS model has incorporated a mass fixer which is comparable to the non-iterative global filter; however, it is not using shape-preserving limits.

4.4 Global filter test

To illustrate the effect of a global filter compared to a local filter, a simple one-dimensional advection test has been constructed. The test considers a classical cyclic rotation test

with a step function as well as a smooth cosine function. The step function will generate large under- and overshoots when advected, whereas the cosine function will be well-represented under advection, thus resulting in an asymmetric distribution of oscillations. The setup is defined as

$$\psi_i = \begin{cases} 0.9 & , \text{ for } 15 < i < 28 \\ 0.5 + 0.4 \cdot \cos\left(\frac{-2\pi(i-60)}{24}\right) & , \text{ for } 60 < i < 84 \\ 0.1 & \text{ otherwise,} \end{cases} \quad (16)$$

where $i = 1 \dots 100$. The above function has equal mass in the two halves of the domain.

The error measures used in Table 1 are the following three commonly used norms:

$$l_1 = \frac{\sum_{k=1}^K |\phi_k - \psi_k|}{\sum_{k=1}^K |\psi_k|}, \quad (17)$$

$$l_2 = \frac{\sqrt{\sum_{k=1}^K (\phi_k - \psi_k)^2}}{\sqrt{\sum_{k=1}^K (\psi_k)^2}}, \quad (18)$$

$$l_\infty = \frac{\max[|\phi_k - \psi_k|]}{\max[|\psi_k|]}, \quad (19)$$

where ϕ_k is the forecast value and ψ_k is the analytic or true value. $\Delta_{\text{mass}}^{\text{step}}$ and $\Delta_{\text{mass}}^{\text{cosine}}$ are the mass changes in the left and right half of the domain containing the step function and the cosine hill, respectively. ϕ_{max} and ϕ_{min} are the final maximum and minimum values in the domain.

As shown in Fig. 3 and Table 1, the unfiltered solution is behaving as expected, generating both undershoots and overshoots with ϕ_{max} up to 0.980 and ϕ_{min} down to 0.059. The filtered solutions are both shape preserving, but the global filter leads to a dampening of the maximum value to 0.883 after two revolutions and 0.853 after three revolutions, whereas the ILMC filter still retains the correct extrema. The filters ability to conserve the mass of each shape, i.e. locality, is also different. The unfiltered solution has a small redistribution caused by the cubic interpolation scheme which, although local, will “spread” every time step. The globally filtered solution has a larger mass difference caused by the scheme’s global nature. With every over- and undershoot, which are largest around the steep gradients, the mass is distributed in the entire domain from areas with small gradients to areas with steeper gradients. The ILMC filter does not have any mass change in the two halves of the domain. Looking at the error norms the unfiltered solution consistently has lower error values than the globally filtered solution. This is caused by the fact that even though the globally filtered solution describes the step function well, the mass redistribution causes the solution of the cosine hill to be degraded, whereas the unfiltered solution describes it accurately. The ILMC filter has consistently the lowest error norms although it is dampening the solution of both the cosine hill and the step function.

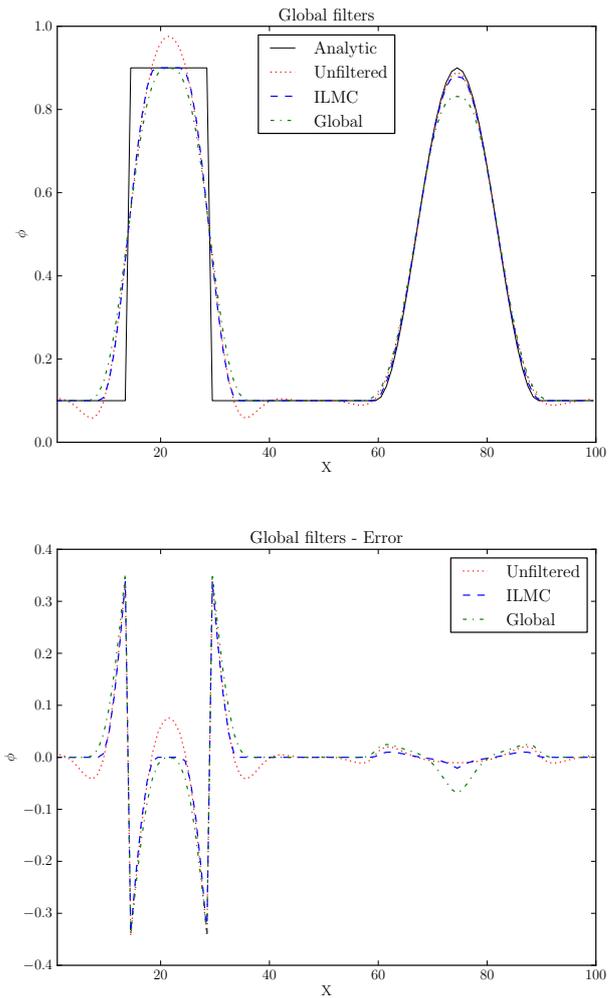


Fig. 3. Filtering: mass-conserving and shape-preserving global filter compared with mass-conserving and shape-preserving local filter. The top panel shows the solution after two revolutions, where the solid line (black) is the analytical solution, the dotted (red) line is the unfiltered LMCSL solution, the dashed (blue) line is the local ILMC filter, and the dot-dashed (green) line is the global shape-preserving filter. The bottom panel shows the difference between the analytical solution and the individual simulations.

5 Results

5.1 NWP comparison

The new scheme is applied for the transport of water in the atmosphere, i.e. water vapour, cloud water, and cloud ice. The HIRLAM model, on which Enviro-HIRLAM is constructed, is a well-proven NWP model and it is therefore important that the LMCSL-3D transport scheme does not degrade the forecast. In Fig. 4 the LMCSL-3D scheme has been used without chemistry for 3 day hindcast, from 1 July 2009, 00:00 UTC, and compared to the HIRLAM model. The hindcasts are not compared to observations as the purpose is to

Table 1. Error norms for global filter test: the table shows the three solutions' error norms as well as mass change for the step and cosine hill functions, and the maximum and minimum values. The velocity is 0.5 and the domain is 100 units long, so the time steps correspond to one, two, and three revolutions.

Filter	Time steps	l_1	l_2	l_∞	$\Delta_{\text{mass}}^{\text{step}}$	$\Delta_{\text{mass}}^{\text{cosine}}$	ϕ_{max}	ϕ_{min}
Unfiltered	200	1.05×10^{-1}	1.83×10^{-1}	3.61×10^{-1}	-9.3×10^{-5}	9.3×10^{-5}	0.966	0.060
Global-M	200	1.09×10^{-1}	1.90×10^{-1}	3.73×10^{-1}	5.7×10^{-2}	-5.7×10^{-2}	0.900	0.100
ILMC	200	7.99×10^{-2}	1.71	3.53	0.0	0.0	0.900	0.100
Unfiltered	400	1.27×10^{-1}	1.95×10^{-1}	3.76×10^{-1}	9.9×10^{-4}	-9.9×10^{-4}	0.980	0.060
Global-M	400	1.41×10^{-1}	2.12×10^{-1}	3.90×10^{-1}	8.8×10^{-2}	-8.8×10^{-2}	0.883	0.100
ILMC	400	9.94×10^{-2}	1.88	3.68	0.0	0.0	0.900	0.100
Unfiltered	600	1.40×10^{-1}	2.03×10^{-1}	3.80×10^{-1}	2.6×10^{-3}	-2.6×10^{-3}	0.971	0.059
Global-M	600	1.69×10^{-1}	2.29×10^{-1}	4.03×10^{-1}	9.1×10^{-2}	-9.1×10^{-2}	0.853	0.100
ILMC	600	1.14×10^{-1}	1.99×10^{-1}	3.75×10^{-1}	0.0	0.0	0.900	0.100

assess the impact it has compared to the default scheme. An in-depth validation is not in the scope of this paper. The left column shows the mean sea level pressure (MSLP) as contour lines and the 850 mb temperature in coloured contours. The right column shows the 6 h accumulated precipitation field from 66 h to 72 h. The top row shows the traditional SL scheme and the bottom row shows the LMCSL-3D scheme, respectively. It is clear that the two hindcasts are similar and difficult to distinguish and that the new transport scheme does not alter the forecast significantly in this case. If one looks closely at the precipitation plots, some differences are present. The general pattern remains the same, but the small areas with strong convective precipitation do change somewhat, both in strength and location. This is to be expected since small changes in the water distribution in the atmosphere will change where the precipitation is triggered. Overall the new scheme seems to provide reasonable hindcasts without the need for additional adjustments and re-tuning of the model.

5.2 Point sources

The main contribution to mass gain using the traditional SL scheme occurs when a very low (possibly zero) background concentration is combined with point sources. This situation generates a considerable amount of negative undershooting which, if not filtered properly, lead to a mass gain when truncated to ensure positive definiteness. To test the mass-conserving properties of the new scheme, a simple worst case scenario, similar to the one used by Chenevez et al. (2004), has been constructed. A point source, with zero background concentration, is emitted in the centre of the domain at the surface, and all deposition has been disabled. The emission consists of 1000 kg of passive tracer emitted linearly during 10 time steps (1 h). As shown in Fig. 5 the conventional SL scheme, with truncation of negative values, increases mass by more than 75 % within a period of 15 h, while the LMCSL-3D scheme conserves mass to machine accuracy.

The mass is calculated as total mass in the entire domain, and in the duration of the simulation no mass escapes the domain. The meteorological conditions are the same as in Sect. 5.1, i.e. starting from 1 July 2009, 00:00 UTC.

5.3 Emissions and large Courant numbers

Semi-Lagrangian schemes are not only efficient because they can be highly multi-tracer efficient but also due to the stability of the method at large Courant numbers (CN). However, it has been suggested that transport schemes cannot use this advantage of stability for very large time steps when used in an ACT model (Hansen et al., 2011). This is caused by inaccuracy in the emissions if the trajectory is longer than one grid box distance, thus “skipping” emissions along the trajectory.

Since Enviro-HIRLAM is semi-Lagrangian and can run with large Courant numbers, it has been tested and found not to introduce any significant errors if the emissions are handled properly and the stencil of the interpolation is large enough. The emission is divided into two halves – one part before advection and the other after the advection. The first half of the emission is thus included in the diffusion and advection, effectively distributing it along the trajectory³. The second half is emitted after advection but before the tendencies from the physics and chemistry modules are calculated, thus representing the latter part of the time step which is not advected along the full trajectory. To simulate the optimal emissions, the model has been run with a small time step (5 min), corresponding to an average CN of 0.3 with a maximum of 1.3, to create a benchmark emission scenario. The second run has a time step of 20 min, corresponding to an average CN of 1.2 with a maximum of 5.0. The maximum CN is a hard limit for Eulerian type schemes, which are unstable for $\text{CN} > 1$, thus posing a constraint for the general time step, but for SL schemes this is not the case and we can choose the time step length more freely. The CN that sets the

³The order of interpolation combined with the horizontal diffusion must be sufficient to cover the entire trajectory.

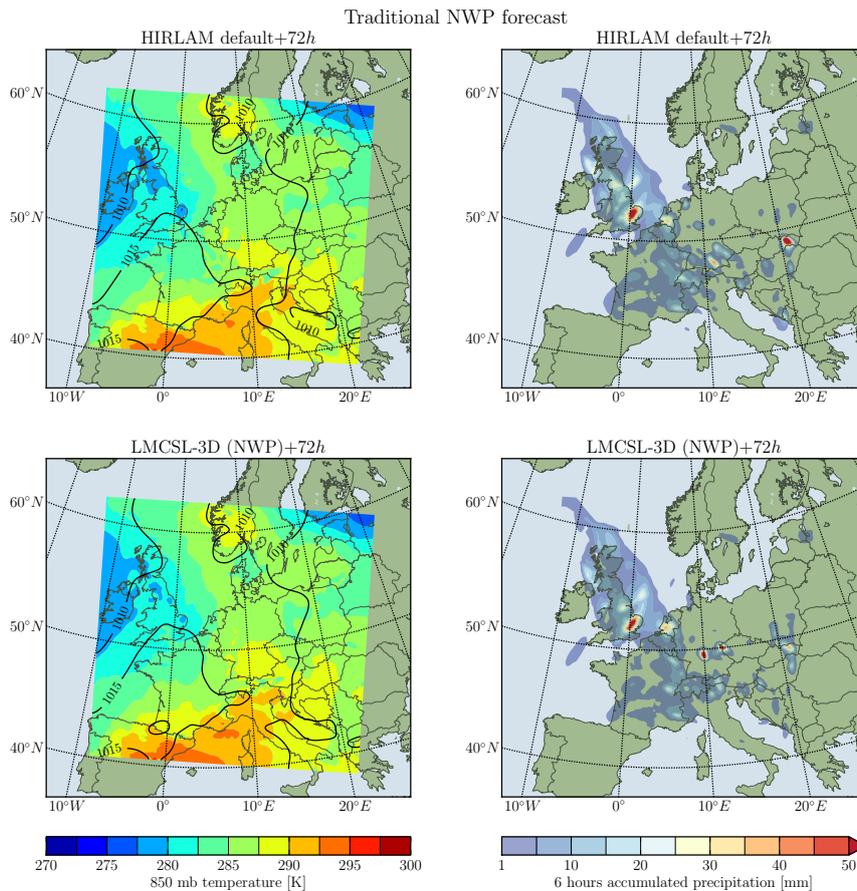


Fig. 4. Comparison between HIRLAM and LMCSL-3D in a 72 h NWP forecast. Upper row shows the traditional HIRLAM forecast, and in the lower row the corresponding LMCSL-3D forecast is shown. The left column shows the mean sea level pressure as contour lines on top of an 850 mb temperature contour field. The right column shows 6 h accumulated total precipitation.

limit for the emission is, however, not the global maximum, but the maximum close to the surface (for surface emissions), which in this case is chosen to be the four lowest levels since we are using cubic interpolations. The relevant CN for the emission is thus: average CN of 0.2 and maximum CN 0.5 for the benchmark run and average CN of 0.8 and maximum CN of 1.8 for the test run, meaning that we are well within the acceptable range.

The test scenario used here is the first European Tracer Experiment (ETEX-1) (Graziani et al., 1998; Nodop et al., 1998), which was carried out in October 1994, in which a non-depositing trace gas was emitted in western France under controlled conditions and then observed across Europe for the next three days. It presents the model with all the above-mentioned challenges and provides observations for comparison. In Fig. 6 it is evident that the LMCSL-3D method performs well and gives a smooth and realistic solution, even with large time steps. Looking at the two simulations, in particular for the first part of the simulation, some differences can be seen in the shape of the tracer cloud; this is mainly caused by the decrease in horizontal diffusion since

the number of time steps are reduced by a factor of four. Another reason is that the trajectories are less accurate the longer they become, and thus small-scale features are not represented to the same degree. Without going into a detailed analysis on the individual observation stations, it can still be concluded that both simulations provide a similar and realistic dispersion of the tracer when compared to the measurements.

In Fig. 7 the LMCSL-3D solution is compared with the non-conserving SL scheme. The upper four panels show the two numerical schemes – SL to the left and LMCSL-3D to the right, 24 h and 48 h after the tracer release. It is clearly evident that the schemes are very similar, but it is possible, in particular after 48 h, to see differences at the plume edge, where the non-conserving scheme seems to have a slightly larger spread, and in the middle of the plume the non-conserving scheme shows a small increase in concentration. In the lower two panels the maximum tracer concentrations during the two simulations are plotted against each other. The left panel shows the first 40 h (when the emission takes place) and the right plot shows the last 40 h, when the plume is

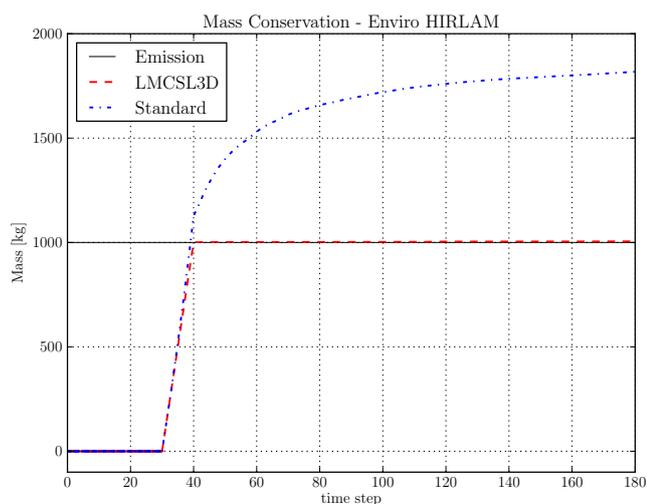


Fig. 5. Mass conservation. The traditional SL scheme shown by the dot-dashed (blue) line and the LMCSL-3D scheme shown by the dashed (red) line. The horizontal line indicates the true emitted amount.

transported passively. An effect of the large time step can be seen in the left panel, where the maximum concentration during the emission is very large. This is because the time step is 20 min and the tracer gas is emitted in a single grid box. The non-conserving scheme only generates a small increase in the maximum value compared with the LMCSL-3D scheme; this is expected since the mass gain is at the plume edge where the negative values are generated. In the right panel the same is evident, but we can see that after the emission has ended the maximum values show a much more realistic level. The differences in maximum concentration are mainly due to the lack of shape preservation in the non-conserving scheme and not because of general mass gain.

5.4 Filter choice

Since the main computational cost of the new advection scheme is the mass-conserving filter and not the advection itself, different filters have been tested. The filters are (see Sect. 4) the ILMC-filter, two global filters, Global-PD (positive definite) and Global-M (monotone), and the DEPDEP benchmark filter. All filters conserve mass to machine accuracy but have different impacts on the chemistry.

Two fields are compared: the ozone (O_3) field and that of sulfur dioxide (SO_2). These two species represents different types of reactions. O_3 is a secondary species – that is, there are no emissions, only initial conditions, boundaries, and chemical reactions – while SO_2 is an almost passive species, with only one chemical reaction (with hydroxide to form sulfate and water) in the condensed CBM-Z scheme and strongly determined by emissions and deposition.

In Fig. 8 the 72 h hindcast using the condensed CBM-Z chemical scheme is shown. The plot shows the LMCSL-3D

scheme using the DEPDEP filter (the reference solution) in the top row. In the second row the truncated SL simulation is shown, and in the third row the ILMC simulation is shown. The fourth row shows the Global-M solution and the bottom row shows the Global-PD solution.

In Fig. 9 a series of comparison plots are shown. The upper two panels show the total mass in lowest model (from the surface and up to approximately 30 m) layer of the two species. O_3 is surprisingly similar although the two global solutions are generally lower than the other solutions, and in particular the positive definite is consistently lower. For SO_2 all solutions are similar, only the truncated solution (not mass conserving) is consistently higher. The SO_2 field is controlled by emissions and deposition, and it is not unexpected that even the non-conserving scheme does not drift from the other filters. For O_3 it seems that although the SL scheme is not conserving, the balance between the individual chemical species and the daily cycle is enough to constrain the solution. The second row shows the average difference in volume density (lowest model layer) between the DEPDEP-filtered solution and the other filters. In both panels the ILMC-filtered solution is closest to the DEPDEP solution. Here the same pattern as before is apparent, the non-conserving scheme is not worse than the conserving schemes; however, the global positive definite filter, in particular for O_3 , has consistently larger differences than the other solutions. It should be noted that it seems that the difference for the positive definite filter has stopped increasing, whereas the differences for the other solutions are still increasing. It will require longer simulations to investigate whether the average difference reaches a constant level or whether the positive definite solution starts increasing again. There is, however, no reason to expect that the average difference should stop increasing, since the simulations will continue to diverge with time. In the lower two rows the largest differences between individual cells are shown: overshoots compared to DEPDEP in the third row and undershoots in the bottom row. It can be seen that for O_3 only the global-positive definite filter is much different; it has differences of the same magnitude as the field itself (see Fig. 8). In the SO_2 case the two positive definite solutions (truncated and Global-PD) are generally similar, with large differences in density. The same large differences can be seen for the global monotone filter, but only for the negative differences, indicating that at some overshoots, mass is distributed elsewhere or, in the positive definite cases, simply not treated.

5.5 Efficiency

The new scheme using the default ILMC filter is compared to the traditional SL transport scheme in two different model setups. In the first setup it is used as a conventional NWP model, i.e. excluding chemistry, with only four advected tracers (specific humidity, cloud water, cloud ice, and TKE). The second setup is the Enviro-HIRLAM setup which includes

ETEX - dependency on time step length

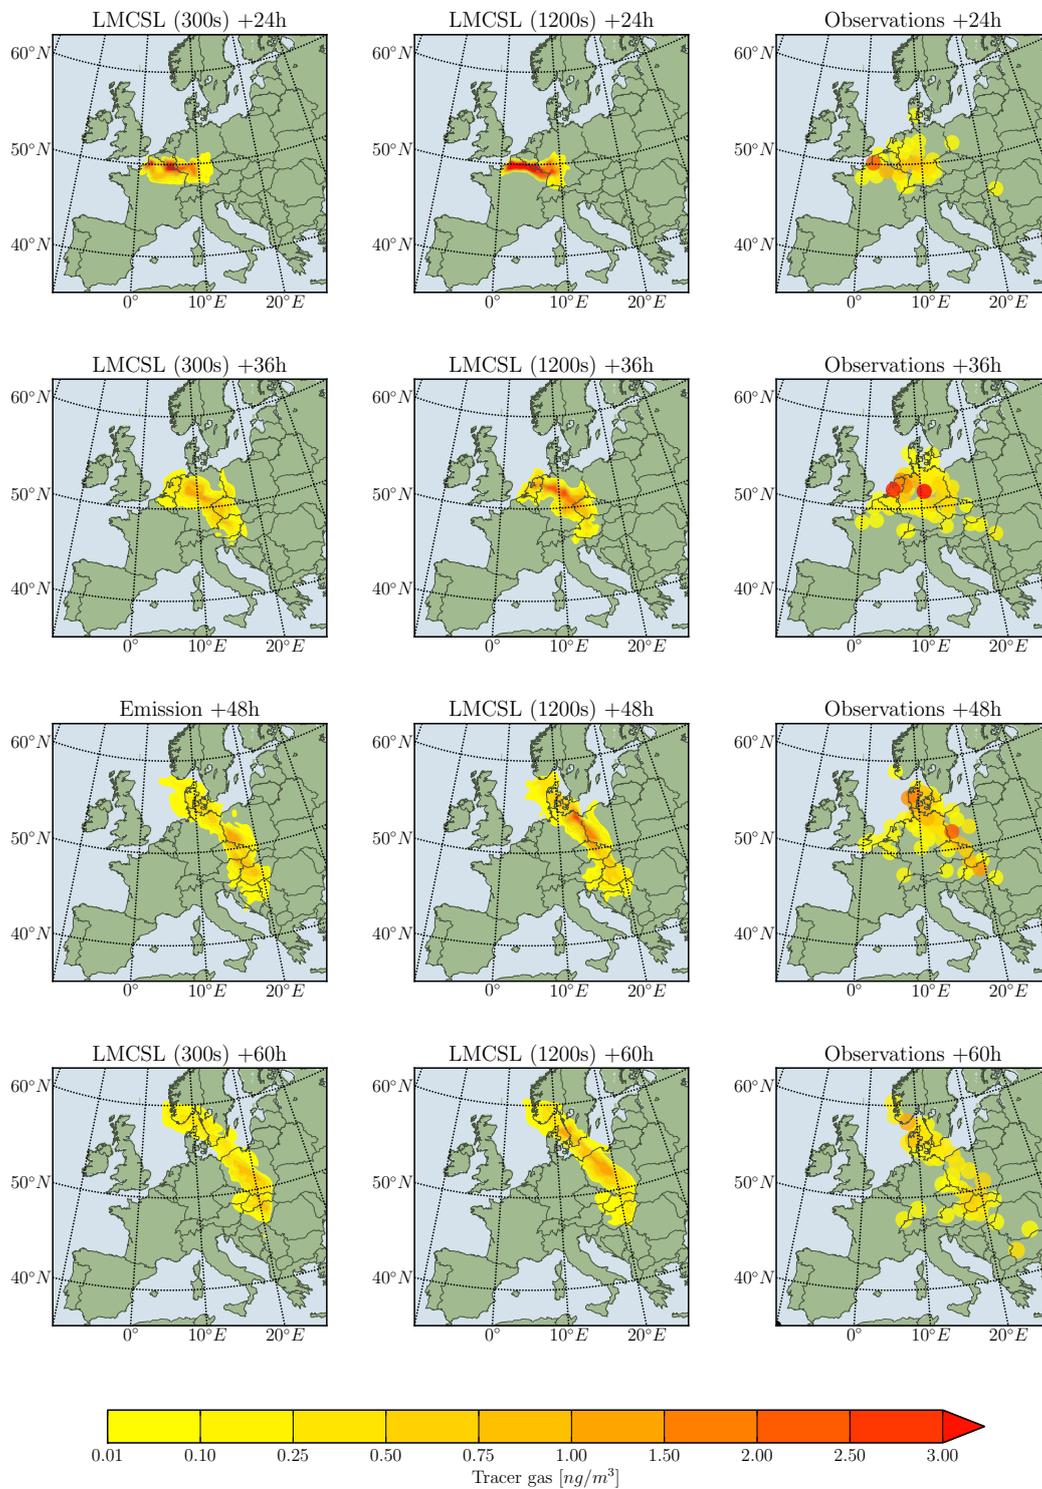


Fig. 6. Emissions with large Courant numbers using the ETEX-1 campaign. The left column shows the simulation using a short time step of 300 s ($\text{CN} \leq 1.3$), the centre column shows the simulation using a long time step of 1200 s ($\text{CN} \leq 5.0$), and the right column shows the observations from the original experiment.

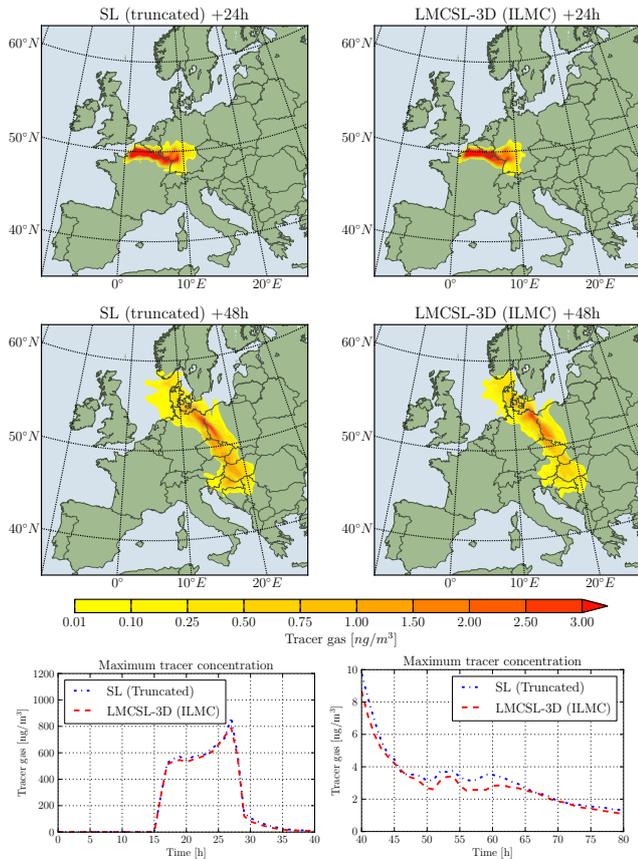


Fig. 7. Point source emission comparison using the ETEX-1 campaign. The top four panels show the truncated SL emission on the left and the LMCSL-3D with ILMC filter on the right. The lower panels show the maximum concentration during the simulations, with the dot-dashed (blue) line being the traditional non-conserving SL scheme and the dashed (red) line is the LMCSL-3D scheme with the ILMC filter.

the condensed CBM-Z chemical scheme and the advection of an additional 32 chemical species. In both setups the meteorological time step is 600 s and the chemistry sub-steps with a time step of 300 s for the Enviro-HIRLAM setup. The number of grid cells is $154 \times 148 \times 40$ and the number of processors is 120. Table 2 shows that in the traditional HIRLAM NWP setup, the new method is approximately 9 % slower and that in the full Enviro-HIRLAM setup the additional cost drops to less than 1 %. The difference is caused by three elements: the additional computations required to calculate the LMCSL interpolation weights, which are independent of the number of variables; the ILMC filter computations, which are proportional to the number of advected tracers; and the chemical scheme itself, which is computationally expensive. This means that, particularly when including chemistry, the additional cost of the new scheme is negligible.

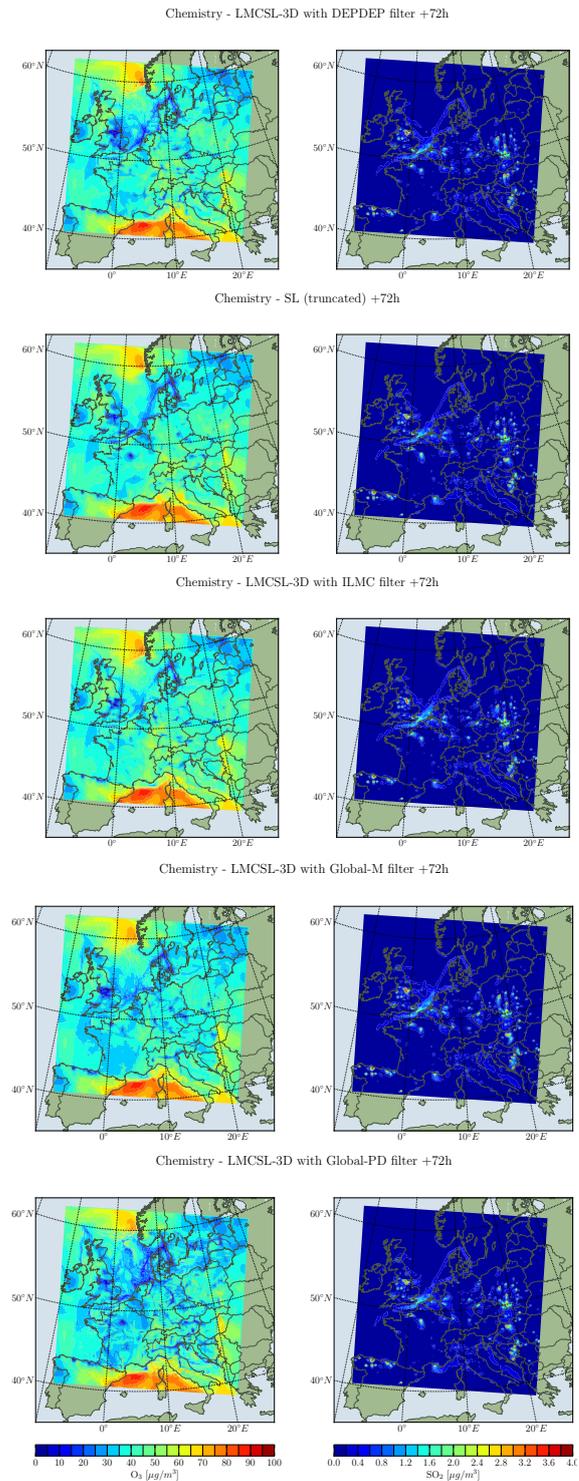


Fig. 8. Chemistry forecast. The panels on the left show the ozone field and the panels on the right show the sulfur dioxide field after 72 h. The top row is the DEPDEP-filtered solution, the second row is the truncated SL, the third row is the ILMC-filtered solution, the fourth row is the globally shape preserving, and the bottom row is the globally positive definite solution. The simulation is from 1 July 2009, 00:00 UTC, to 4 July 2009, 00:00 UTC.

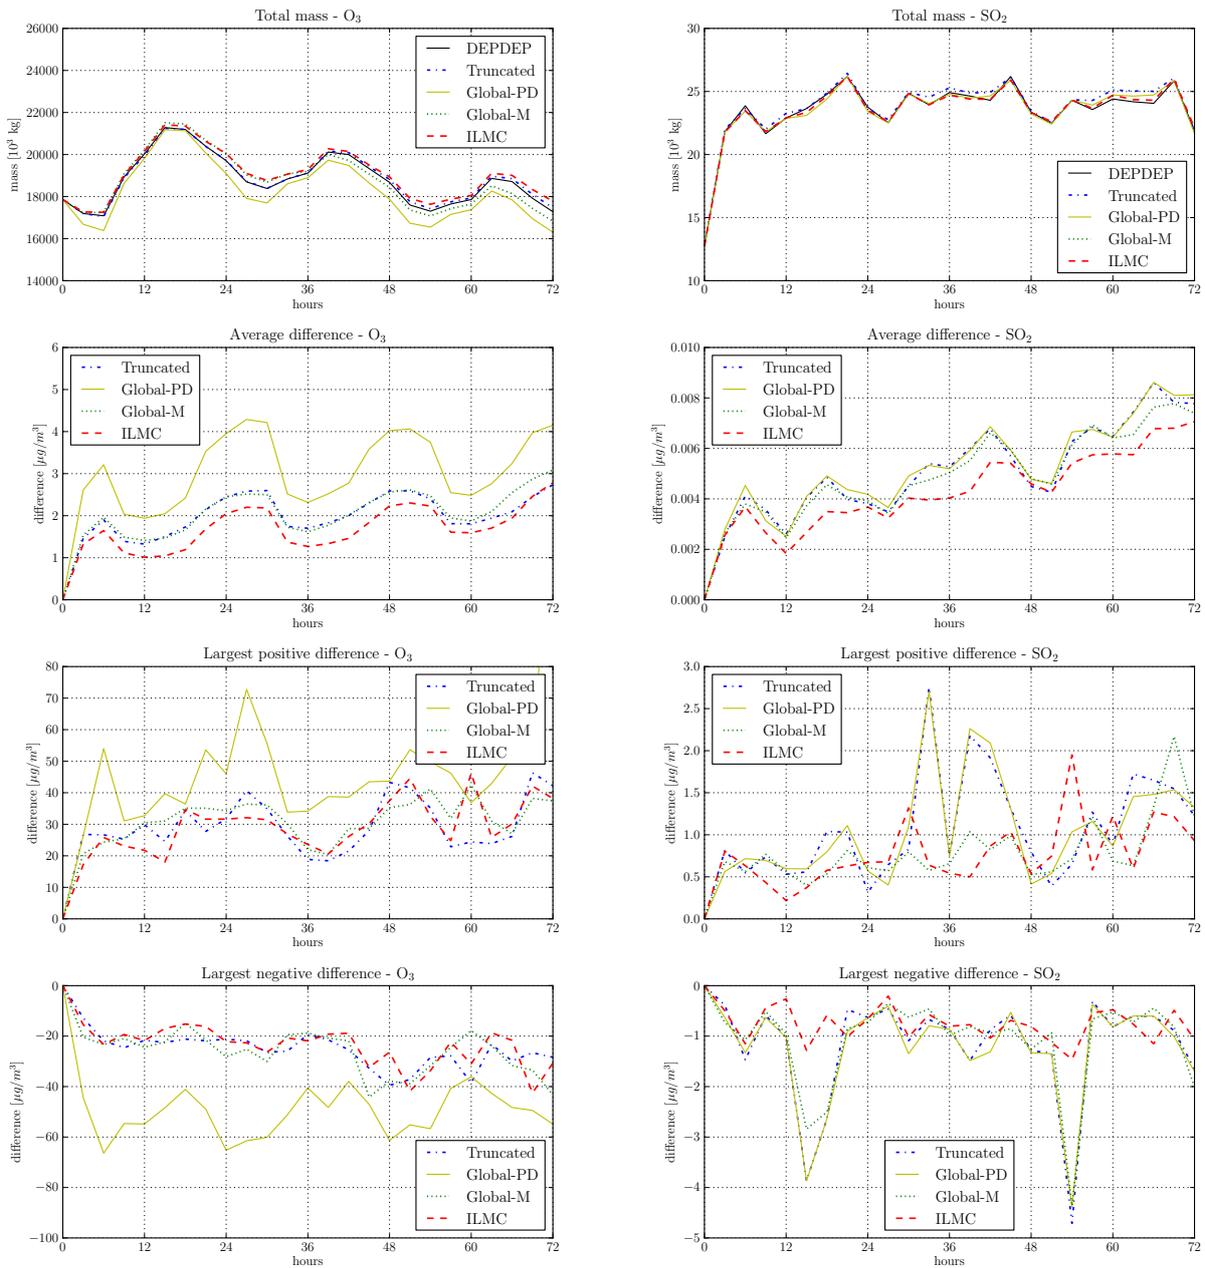


Fig. 9. O₃ (left column) and SO₂ (right column) forecast from 1 July 2009, 00:00 UTC, to 4 July 2009, 00:00 UTC (see Fig. 8). Top row: total mass in the surface layer. Second row: average volume density difference compared to the DEPDEP-filtered solution. Third row: largest positive difference compared to the DEPDEP-filtered solution. Bottom row: largest negative difference compared to the DEPDEP-filtered solution. The lines are as follows: thin solid line (black) is the DEPDEP solution, dot-dashed line (blue) is the truncated SL solution, thick solid line (yellow) is the global positive-definite-filtered solution (Bartnicki), dotted line (green) is the global monotone solution, and the dashed line (red) is the ILMC-filtered solution.

5.6 Vertical mixing

When including chemical constituents the amount of numerical vertical mixing is important since chemical reactions are triggered if the species are mixed unphysically. In Fig. 10 a measure of the vertical mixing is shown for the traditional

SL scheme and the LMCSL-3D scheme with both the ILMC filter and the DEPDEP filter. The amount of numerical vertical mixing should not be increased or changed significantly using the new scheme. To estimate the mixing, each layer in the model is initialized with 100 kg of passive tracer in each grid box. The layers are treated separately, i.e. each layer is

Table 2. Efficiency comparison between the SL scheme and the LMCSL-3D scheme when used as traditional NWP model and when using an online integrated condensed CBM-Z chemistry scheme with 32 additional advected tracers. The time in column two and three are the average time per time step during a 24 h forecast.

Simulation	SL	LMCSL-3D	Efficiency
HIRLAM	0.428 s	0.470 s	91.2 %
E-HIRLAM	17.016 s	17.116 s	99.4 %

an independent passive tracer, which does not mix with the other tracers. It is then possible to see to what extent the emitted mass is preserved in the emission layer. The simulation used here is the same synoptic situation as in Sect. 5.1 and the length of the simulation is 24 h. In Fig. 10 the percentage of mass which is vertically mixed, i.e. lost from the emission layer, is shown⁴. The three schemes have almost identical vertical mixing profiles, which indicates that although the LMCSL-3D schemes do alter the interpolations slightly and the filters redistribute some of the mass, the general degree of mixing is almost unchanged. In the upper five layers (10 hPa to 100 hPa) the LMCSL-3D schemes show a slight increase in mixing, which could be an effect of this particular meteorological situation; however, it is more likely due to less vertical motion (indicated by less mixing), which will tend to increase the generation of undershoots due to larger gradients between the layers, thus activating the filter more often. The mixing in the DEPDEP-filtered solution is slightly less than for the ILMC filter. This indicates that for this particular forecast it is impossible to decrease the mixing more without sacrificing mass conservation.

6 Conclusions

In this paper a new transport scheme is presented. The new scheme, LMCSL-3D with ILMC filter, which is mass conserving, shape preserving, and multi-tracer efficient, has been implemented in the online integrated NWP/ACT Enviro-HIRLAM model. The new scheme is shown to be almost as efficient as the original scheme – that is, the additional computational cost is less than 1 % – while improving on several of the so-called desirable properties, in particular mass conservation. When used as a traditional NWP model without chemistry, i.e. the HIRLAM setup, the scheme gives nearly identical results to that of HIRLAM itself for the cases considered here. When run with full chemistry, the results demonstrate that, even for short lead times, the mass-conserving and shape-preserving properties of a numerical scheme affect the chemistry.

⁴The data have been normalized since the SL scheme increases total mass. In the actual test, mass was approximately 25 % larger than in the LMCSL-3D schemes (with the ILMC and DEPDEP filter).

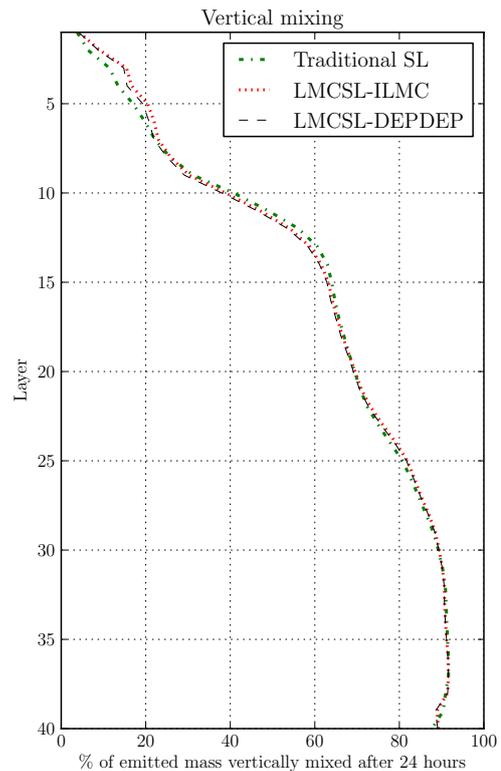


Fig. 10. Vertical mixing of passive tracers. This figure shows the amount of mass which has been mixed vertically, i.e. not present in the layer in which it was emitted. The dot-dashed (green) line is the traditional non-conserving SL scheme, the dotted (red) line is the LMCSL-3D scheme with the ILMC filter, and the dashed (black) line is the LMCSL-3D scheme with the DEPDEP filter.

It is also shown that the mass conservation property is not enough to guarantee realistic forecasts; locality and shape preservation are just as important. Global filters redistribute mass remotely and can potentially degrade the details to a degree where key effects disappear. Longer simulations and comparison with observations are required to appreciate if any of the methods will diverge towards unphysical solutions.

The last point addressed in this paper is the performance of the emission method. It is shown that the semi-Lagrangian scheme can have accurate emissions, even when using long time steps (Courant numbers up to 5).

It is noted that although the LMCSL-3D scheme is using the same interpolation weights to increase consistency, the mass field itself, which is described by the surface pressure and a terrain-following coordinate (σ - η hybrid), is not advected using the LMCSL-3D scheme. Thus the so-called mass wind inconsistency is still present. It will require substantial changes to the model to formally remove this inconsistency, if at all possible.

Acknowledgements. This work is done as part of the Centre for Energy Environment and Health (CEEH), financed by The Danish Strategic Research Programme on Sustainable Energy under contract no. 09-061417.

Edited by: O. Marti

References

- Baklanov, A.: Integrated Meteorological and Atmospheric Chemical Transport Modeling: Perspectives and Strategy for HIRLAM/HARMONIE, HIRLAM newsletter 53, 68–78, 2008.
- Baklanov, A.: Chemical weather forecasting: a new concept of integrated modelling, *Adv. Sci. Res.*, 4, 23–27, doi:10.5194/asr-4-23-2010, 2010.
- Baklanov, A., Mahura, A., and Sokhi, R. (Eds.): *Integrated Systems of Meso-Meteorological and Chemical Transport Models*, Springer, XVI, ISBN: 978-3-642-13980-2, 2011.
- Bartnicki, J.: Simple Filtering Procedure for Removing Negative Values from Numerical Solutions of the Advection Equation, *Environ. Softw.*, 4, 187–201, 1989
- Bott, A.: A Positive Definite Advection Scheme Obtained by Non-linear Renormalization of the Advective Fluxes, *Mon. Weather Rev.*, 117, 1006–1015, 1989.
- Chenevez, A., Baklanov, A., and Sørensen, J. H.: Pollutant transport schemes integrated in a numerical weather prediction model: model description and verification results, *Meteorol. Soc.*, 11, 265–275, 2004.
- Flemming, J., Inness, A., Flentje, H., Huijnen, V., Moinat, P., Schultz, M. G., and Stein, O.: Coupling global chemistry transport models to ECMWF's integrated forecast system, *Geosci. Model Dev.*, 2, 253–265, doi:10.5194/gmd-2-253-2009, 2009.
- Flemming, J., Huijnen, V., Jones, L., ArtetaFrance, J., and Masart, S.: On-line integration of atmospheric chemistry in the Integrated Forecasting System, Macc, technical documentation of the c-ifs, MACC, Monitoring Atmospheric Composition and Climate, available at: <http://www.gmes-atmosphere.eu/documents/deliverables/g-rg/> (last access: 17 July 2013), 2012.
- Graziani, G., Klug, W., and Mosca, S.: Real-Time Long-Range Dispersion Model Evaluation of the ETEX First Release, EUR 17754 EN. Office for official publications of the European Communities, Luxembourg, ISBN 92-828-3657-6, 1998.
- Hansen, A. B., Brandt, J., Christensen, J. H., and Kaas, E.: Semi-Lagrangian methods in air pollution models, *Geosci. Model Dev.*, 4, 511–541, doi:10.5194/gmd-4-511-2011, 2011.
- Jockel, P., von Kuhlmann, R., Lawrence, M., Steil, B., Breninkmeijer, C., Crutzen, P., Rasch, P., and Eaton, B.: On a fundamental problem in implementing flux-form advection schemes for tracer transport in 3-dimensional general circulation and chemistry transport models, *Q. J.-Roy. Meteorol. Soc.*, 127, 1035–1052, 2001.
- Kaas, E.: A simple and efficient locally mass conserving semi-Lagrangian transport scheme, *Tellus*, 60, 305–320, 2008.
- Kaas, E. and Nielsen, J. R.: A Mass-Conserving Quasi-Monotonic Filter for Use in Semi-Lagrangian Models, *Mon. Weather Rev.*, 138, 1858–1876, 2010.
- Korsholm, U. S.: Integrated modeling of aerosol indirect effects, DMI scientific report, DMI, 1, available at: <http://www.dmi.dk/dmi/print/sr09-01.pdf> (last access: 17 July 2013), 2009.
- Korsholm, U. S., Baklanov, A., Gross, A., Mahura, A., Sass, B. H., and Kaas, E.: Online coupled weather forecasting based on HIRLAM – overview and perspective of Enviro-HIRLAM, HIRLAM newsletter 54, 151–167, 2008.
- Lauritzen, P. H., Kaas, E., and Machenhauer, B.: A Mass-Conservative Semi-Implicit Semi-Lagrangian Limited-Area Shallow-water Model on the Sphere, *Mon. Weather Rev.*, 1205–1221, 2006.
- Lauritzen, P. H., Kaas, E., Machenhauer, B., and Lindberg, K.: A mass-conservative version of the semi-implicit semi-Lagrangian HIRLAM, *Q. J. Roy. Meteorol. Soc.*, 134, 1583–1595, 2008.
- Lawrence, M. G., Hov, Å., Beekmann, M., Brandt, J., Elbern, H., Eskes, H., Feichter, H., and Takigawa, M.: The Chemical Weather, *Environ. Chem.*, 2, 6–8, 2005.
- Machenhauer, B., Kaas, E., and Lauritzen, P.: Finite-Volume Methods in Meteorology, Computational Methods for the Atmosphere and the Oceans: Special Vol. 14, Elsevier, ISBN-13: 978-0-444-51893-4, 2009.
- Nodop, K., Connolly, R., and Girardi, R.: The field campaign of the European Tracer Experiment (ETEX): overview and results, *Atmos. Environ.*, 32, 4095–4108, 1998.
- Rasch, P. J. and Williamson, D. L.: Computational aspects of moisture transport in global models of the atmosphere, *Q. J. R. Meteorol. Soc.*, 116, 1071–1090, 1990.
- Sørensen, B., Kaas, E., and Lauritzen, P.: A Locally Mass-Conserving semi-Lagrangian Regional NWP model with conserved quasi-Lagrangian vertical coordinates, *Q. J. R. Meteorol. Soc.*, in review, 2012.
- Thuburn, J.: Some conservation issues for the dynamical cores of NWP and climate models, *J. Comput. Phys.*, 227, 3715–3730, 2008.
- Thuburn, J. and McIntyre, M.: Numerical advection schemes, cross-isentropic random walks, and correlations between chemical species, *J. Geophys. Res.*, 102, 6775–6797, 1997.
- Zaveri, R. A. and Peters, L. K.: A new lumped structure photochemical mechanism for large-scale applications, *J. Geophys. Res.*, 104, 30387–30415, 1999.