

CARBON ATMOSPHERIC TRACER RESEARCH TO IMPROVE NUMERICAL SCHEMES AND EVALUATION



CATRINE

Carbon Atmospheric Tracer
Research to Improve
Numerics and Evaluation

D1.2 Progress report on implementation aspects and preliminary results towards an improved tracer advection scheme

Due date of deliverable	30/06/2025
Submission date	24/06/25
File Name	CATRINE D1-2-V1.0
Work Package /Task	WP1
Organisation Responsible of Deliverable	Météo-France
Author name(s)	Sylvie Malardel, Michail Diamantakis, Claire Laurent, Giovanni Tumolo, and Filip Vana.
Revision number	1
Status	Issued
Dissemination Level / location	Public



Funded by the
European Union

The CATRINE project (grant agreement No 101135000) is funded by the European Union.

Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the Commission. Neither the European Union nor the granting authority can be held responsible for them.

1 Executive Summary

The advection scheme is a numerical method used to solve the Partial Differential Equations (PDEs) that represent the resolved (non-parametrized) component of transport. In CATRINE deliverable D1.1 (Diamantakis et al, 2025), a comprehensive evaluation of the core global CO2MVS model (IFS) advection scheme was conducted, focusing on its mass conservation properties and its impact on the accuracy of tracer transport. Here, building on the analysis, findings and recommendations of the D1.1 report, we explore several ideas to improve the tracer advection scheme conservation properties while maintaining a good overall accuracy in tracer transport. These proposed improvements are implemented and tested within the IFS code library, providing insights into their impact on the relevant case studies considered in the D1.1 report.

This work is linked with other work packages—particularly WP5 and WP7—as we leverage protocols developed therein to evaluate the new developments in greenhouse gas simulations under realistic initial and boundary conditions. Detailed comparisons between the different approaches are presented, identifying the most promising options based on mass conservation performance and transport errors relative to observations.

Finally, we provide recommendations for further developments, which will unify and finalise work summarised in this deliverable. This task will take place in CATRINE WP2 providing an improved tracer advection scheme for the global CO2MVS model IFS.

Table of Contents

1	Executive Summary	2
2	Introduction	4
2.1	Background	4
2.2	Scope of this deliverable	4
2.2.1	Objectives of this deliverables	4
2.2.2	Work performed in this deliverable	5
2.2.3	Deviations and counter measures	5
2.3	Project partners:	5
3	Exploring improvements of the COMAD interpolation scheme	6
3.1	Background	6
3.2	COMAD for divergent flows	6
3.3	COMAD along the vertical	7
3.4	More accurate COMAD interpolation weights	8
3.4.1	COMAD 3 points	8
3.4.2	Computation of the COMAD weights “along the SL trajectory”	10
3.5	A robust linear and adjoint model for COMAD	10
4	Investigating the boundary condition of the SL scheme near the surface	11
4.1	Background	11
4.2	Introducing an extra-level below the surface	11
4.3	Modifying COMAD weights near the surface	12
5	Limiters with improved conservation properties for tracers	14
5.1	Impact of limiters in tracer correlation test case	15
5.2	Plume tracer passive advection case study	17
5.3	Testing the limiters on tracers emitted by point sources	18
6	Advection of density versus specific ratio	20
6.1	Background	20
6.2	Results	20
7	Testing the improvements using WP5 and WP7 protocols	22
8	Sweep interpolation: a promising alternative for the semi-Lagrangian scheme for tracer transport	27
9	Concluding remarks and future work	29
10	References	30

2 Introduction

2.1 Background

As summarized in D1.1 CATRINE deliverable, the ECMWF Integrated Forecast System (IFS) is a highly efficient, multi-purpose global model used for both weather forecasting and atmospheric composition. Coupled with a 4-dimensional variational data assimilation system (4D-VAR; Rabier et al., 2000), it delivers accurate atmospheric analyses every six hours, serving as the foundation for global and regional forecasts. It also underpins ERA5, producing widely used reanalysis datasets. Thanks to its advanced numerical accuracy—recognized as the most accurate global model by WMO metrics—and strong capabilities in atmospheric composition, the IFS is a leading choice for 4D-VAR atmospheric inversion systems. Its development benefits from collaborations across European member states, notably with Météo-France and the ACCORD consortium which share the same dynamical core, based on a spectral transform method with a semi-implicit semi-Lagrangian (SISL) time-stepping scheme. This combination allows for long, stable time-steps with minimal phase errors and supports the efficient transport of multiple tracers—an essential feature for atmospheric composition and inversion modelling.

Despite these strengths, the semi-Lagrangian (SL) advection scheme—while accurate, monotonic, and positive-definite—does not inherently conserve mass, except under idealized conditions rarely encountered in practice. As a result, SL advection can introduce global mass conservation errors, particularly relevant in atmospheric composition, climate simulations, and high-resolution forecasts where conservation is critical. To address this, the IFS includes mass fixer algorithms that adjust tracer concentrations post-advection to restore global mass conservation, as developed from early 2D methods (Bermejo and Conde, 2002) to more advanced 3D implementations (Diamantakis and Agusti-Panareda, 2017). These fixers preserve monotonicity and positive-definiteness and have been shown to improve greenhouse gas simulations and moist tracer forecasts. However, because these corrections are applied independently of the governing equations, we would like to reduce the need for such corrections by reducing conservation errors during the advection process itself as much as possible. This has the potential to make the advection scheme and the estimation of fluxes more accurate.

The findings of D1.1, have motivated new research and developments to improve the “inherent” conservation properties of the SL advection scheme of the IFS maintaining its good accuracy. Our progress and the results obtained, which demonstrate this progress, are summarized in the following sections. As is often the case with research endeavours, not all of the proposed ideas have led to significant improvements. However, we have successfully defined a set of changes planned for WP2 that are expected to significantly reduce the mass conservation error produced by the IFS advection scheme. Furthermore, as a result of work in D1.1 and D1.2, we have been able to increase our understanding on the origin and evolution of mass conservation and transport errors in SL advection schemes.

2.2 Scope of this deliverable

2.2.1 Objectives of this deliverables

The objective of this deliverable is to conduct the required research and development towards an accurate advection scheme with improved conservation properties that will result in accurate estimation of emissions. It is the first stage towards the implementation of an improved scheme for CO2MVS which will be completed in WP2. This objective is achieved by (i) exploring monotone limiters with improved conservation (ii) improving the semi-Lagrangian scheme formulations especially the one based on COMAD interpolation (iii) explore

CATRINE

improvements of vertical advection schemes and (iv) improving the stability and robustness of the linear model and adjoint of the COMAD scheme to allow its use in 4D-VAR assimilation and inversion experiments for tracers.

2.2.2 Work performed in this deliverable

The work performed here follows the description of Task 1.2 of WP1 of the CATRINE proposal. More specifically: (i) we have conducted investigations on how to improve the vertical aspects of SL advection which are summarized in section 3.4 and section 4 of this report, (ii) we have explored various avenues for improving the COMAD formulation and its conservative performance which are summarized in section 3.4, (iii) we have produced stable versions of the the COMAD scheme linear model and its adjoint which have been successfully used in inversion experiments (see section 3.5), and (iv) we have developed improved versions of the existing IFS monotone limiters which reduce substantially the mass conservation error of the SL advection scheme (section 5). These new developments have been tested on CATRINE test beds developed in D1.1 while a new test case has been introduced (tracer correlated tests) which has been useful for assessing the accuracy of our improved algorithms. The most promising developments have been assessed in greenhouse gas simulations using the WP5 and WP7 (Chevallier and Agusti-Panareda, 2025) protocols. This testing has helped us to make recommendations for our final developments in WP2. Furthermore, in light of a new promising development which appeared in the scientific literature very recently (see section 8), we have also investigated a new method. This item of work was not described in the original CATRINE proposal but seems to produce very encouraging results.

2.2.3 Deviations and counter measures

None

2.3 Project partners:

Partners	
EUROPEAN CENTRE FOR MEDIUM-RANGE WEATHER FORECASTS	ECMWF
COMMISSARIAT A L ENERGIE ATOMIQUE ET AUX ENERGIES ALTERNATIVES	CEA
METEO-FRANCE	METEO-FRANCE
WAGENINGEN UNIVERSITY	WU
KARLSRUHER INSTITUT FUER TECHNOLOGIE	KIT
HELSINGIN YLIOPISTO	UH
UNIVERSITE DE REIMS CHAMPAGNE-ARDENNE	URCA
ALBERT-LUDWIGS-UNIVERSITAET FREIBURG	UFR

3 Exploring improvements of the COMAD interpolation scheme

3.1 Background

The variant of the SL interpolation weights using the COMAD correction (Malardel and Ricard, 2015) has been operational in the Meteo-France regional model since 2013. It has also been used in the IFS for atmospheric composition tracers since 2021 when IFS model cycle 47R3 became operational. However, the current version of the COMAD correction only applies to horizontal interpolation and it is restricted to horizontally convergent flow.

Also, in its original version, the COMAD interpolation was not coded for the tangent-linear (TL) and adjoint (AD) code of the SL scheme. Consequently, they could not be used in the inversion system based on the IFS 4DVAR scheme.

In section 3, we outline several attempts to improve the conservation properties of the COMAD SL scheme in the non linear model, as well as the developments required to incorporate COMAD interpolation into the 4DVAR. These developments cover the three first points planned in Task 1.2 of CATRINE.

3.2 COMAD for divergent flows

The interpolation weights w' for the COMAD interpolation depend on the wind divergence along the direction of interpolation. In the linear case, they are computed from the traditional weights w as follows:

$w' = \alpha w + 0.5(1 - \alpha)$ where $\alpha = 1 + Ddt$ with D the 1D divergence along the direction of interpolation and dt the time step of the model.

In the current version of the code, α is restricted to values smaller than 1 ($D < 0$), i.e. there is no COMAD correction if the wind along the direction of interpolation is divergent.

A series of tests was run with the α coefficients bounded by 2 instead of 1, as was previously done, in order to include a COMAD correction in the divergent part of flow. These tests were performed for idealised test cases, the WP5/7 intercomparison simulations and weather forecast tests using the COMAD scheme in the transport of specific humidity and cloud fields. All these tests have yielded favourable results.

As an example, we present two cases here: the bubble case and the mountain case with tracer near the surface (see test case descriptions in D1.1). Both show a small improvement for cubic interpolation (figure 3.2a) and for linear interpolation (figure 3.2b) when COMAD is activated also for divergent flows (dotted lines versus dashed lines).

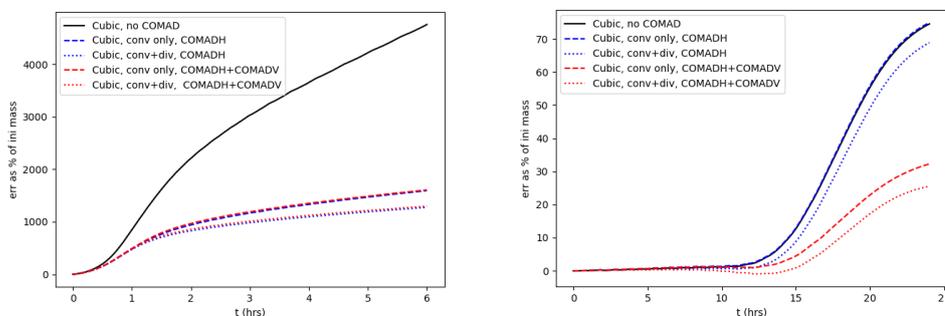


Figure 3.2a. High order cubic interpolation: mass conservation errors (in percent of initial mass) for the bubble test case (left) and the mountain test case with the tracer near the surface (right). See D1.1 for the case descriptions.

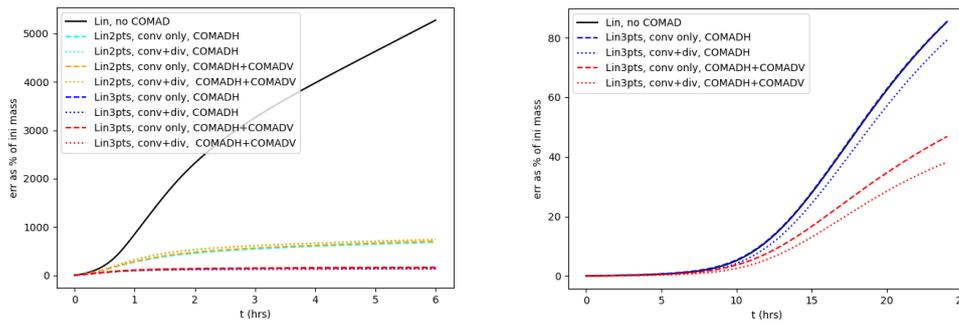


Figure 3.2b. Linear interpolation: mass conservation errors (in percent of initial mass) for the bubble test case (left) and the mountain test case with the tracer near the surface (right). See D1.1 for the case descriptions.

3.3 COMAD along the vertical

In the IFS, the hybrid vertical levels η are based on the hydrostatic pressure. The generalised vertical velocity $\tilde{\eta}$ is used for the SL advection along the vertical.

$D\tilde{\eta}/D\eta$ is a measure of the vertical divergence along the vertical which is consistent with these levels. The SL vertical interpolation weights can then also be modified with the COMAD correction with the same formulation as for the horizontal directions. This option (COMADV) has been tested again with the stress test cases described in D1.1 and then with more realistic meteorology (same as in section 3.2)

In the bubble case, there is a very small negative impact (blue versus red lines on figures 3.2a and 3.2b, left panels, better seen on the zoom, figure 3.4b) but there is a significant positive impact in the mountain case (figures 3.2a and 3.2b, right panels). We’ve also tested this option with a case of intermediate complexity. Unlike the bubble and the mountain cases which are purely adiabatic, the idealised tropical cyclone (TC) test case which has been designed for the model intercomparison exercises DCMIP12 and DCMIP16 (Wilson *et al*, 2024) needs the addition of subgrid parametrizations for cloud and precipitation, vertical subgrid transport and interaction with the surface. For CATRINE, we have used this test case with tracers added in the rain band region where active small scale deep convective clouds develop in the wake of the TC. Figure 3.3 shows the clear positive impact of COMADV in this case. In this plot, we have also included the new 3-point COMAD method which is explained in the following section.

We have also tested COMADV in the greenhouse transport cases presented in section 7 using either WP5 or WP7 emission protocols. For reasons that we do not currently understand we did not see the benefit we see in the idealized case studies here. We find that the mass conservation error increases faster when COMADV is included in the simulation. Adding it seems to increase the mass of greenhouse tracers (CO₂ and CH₄ in particular) mostly in the tropopause zone and in latitudes north of 40°N (not shown here). Overall, results from these more realistic simulations showed that it is better to apply COMAD only in the horizontal.

CATRINE

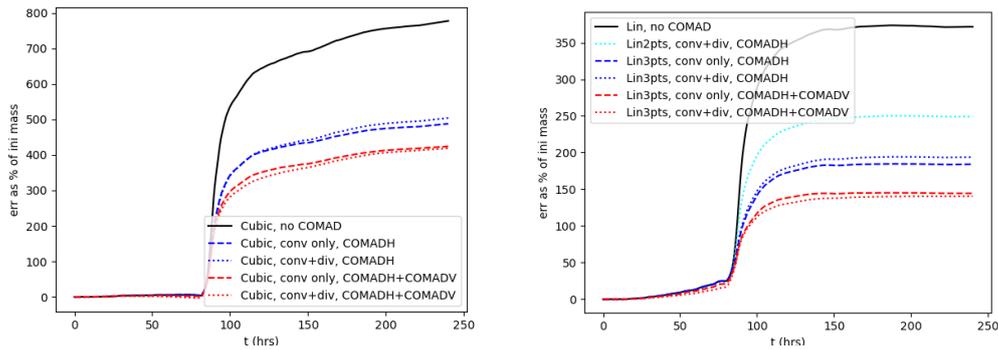


Figure 3.3. DCMIP16 tropical cyclone test case. Tracer mass error for cubic (left) and linear (right) interpolation.

3.4 More accurate COMAD interpolation weights

3.4.1 COMAD 3 points

COMAD originally addressed a problem of the incorrect representation of grid scale horizontal convergence by the SL scheme which triggers very large errors inside the corresponding grid scale (single model column) ascent. The ascending warm bubble case is an archetype of the atmospheric conditions for which SL errors become very large because the grid scale convergence is unperturbed for several hours around a no-wind point.

Traditional SL linear interpolation uses values at the two grid-points around a departure point along the direction of interpolation in order to estimate the value at departure point. For the sake of simplicity, the current linear COMAD interpolation also uses a 2 points stencil. However, analysing the analogy between the linear interpolation and the finite-volume donor scheme shows that the COMAD correction is underestimated by the 2-point formulation in case of grid scale convergence around a no-wind point (bubble case for example). A more accurate 3-point formulation was implemented and tested during CATRINE WP1.

Figure 3.4a (left) shows the tracer distribution in the single column of the ascending buoyant bubble and its horizontal dispersion in the divergent flow where the bubble reaches its level of equilibrium. Because the SL scheme is blind to the grid scale convergence, there is no horizontal transport of clear air at the bottom of the column. In the vertical, the SL scheme transports the high value of tracer originally near the surface within the ascent. As the value at the bottom level remains undiluted by the convergent wind around the no wind point, it behaves as an infinite source of tracer which generates a huge tracer mass conservation error of 5000% after 3 h (figure 3.2b, left).

With COMAD 2 points, the concentration in the wake of the bubble is substantially reduced (figure 3.4a, middle) and the conservation error after 3 h is reduced from 5000 to 700% (figure 3.4b). With the new 3 points formulation of COMAD, the wake of the bubble is almost undetectable (figure 3.4a, right) and the mass increase is now less than 100% of the initial mass (figure 3.4b).

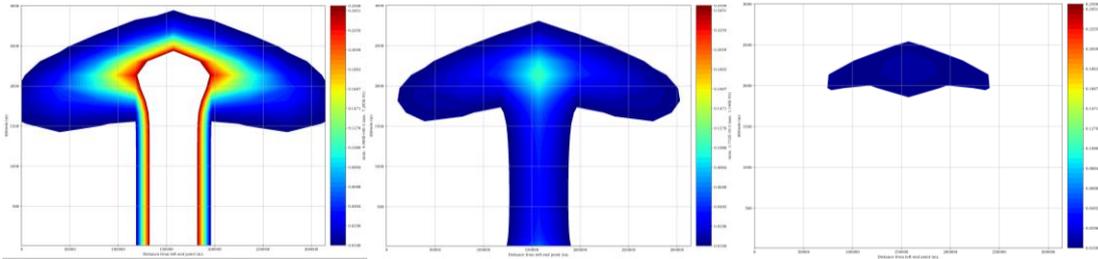


Figure 3.4a. Vertical zonal cross-sections of tracer across the bubble after 3 h of simulation in case of linear interpolation without COMAD (left), with COMAD 2 points (middle) and COMAD 3 points (right). The white colour indicates that the tracer concentration is above the contour plot maximum 0.25. The same contour range is used for all three plots: 0.01-0.25. The initial bubble tracer concentration is 1. The vertical axis shows the height from the surface in meters and the horizontal the distance from the 0-longitude grid-point.

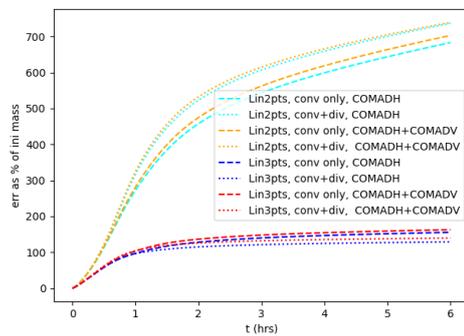


Figure 3.4b. Linear interpolation. Mass conservation errors (percent of initial mass) for the bubble case for COMAD 2 points versus COMAD 3 points.

In the TC case (figure 3.3), there is also a significant gain between COMAD 2 points (cyan) and COMAD 3 points (blue). The generalisation of the more accurate formulation in the case of cubic interpolation is not straightforward because there is no more simple analogy with the finite volume formulation. Improving the accuracy of the higher order interpolation will however still be under investigation during the WP2 of CATRINE.

Finally, the new COMAD 3-point scheme has been tested on the point source emission tests of D1.1 (section 5.2) and has been compared against the standard 2-point scheme. Results from one such experiment are shown in figure 3.4c which shows a very small improvement from the application of the 3-point scheme. Overall, the same behaviour is observed when the tracer point source is placed at a different location than the one used in figure 3.4c.

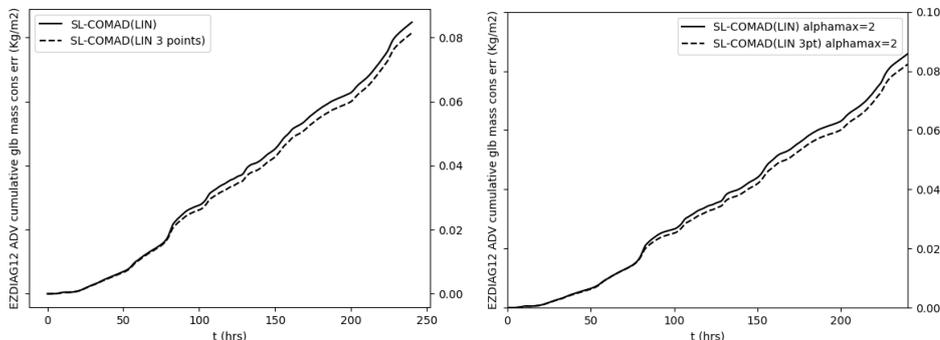


Figure 3.4c: Mass conservation error (in kg/m²) evolution in time (10 days) for China emission point test in winter (see 18 emission point test as described in D1.1 section 5) for COMAD 2 points versus COMAD 3 points in the two cases: COMAD for convergent flows only (alphamax=1) and enhanced COMAD for both convergent and divergent flows (alphamax=2) respectively. This case shows a typical behaviour common to other emission points too.

3.4.2 Computation of the COMAD weights “along the SL trajectory”

In the current code, the α coefficients are computed with the 1D wind divergence of the current time step but at the arrival point (i.e. where the air parcels will be at the next time step). In case of a large CFL¹ number, the position of the departure point can be far from the one of the arrival point. It is then more accurate to compute the alpha coefficient “along the trajectory”. In practice, we implemented a computation using the same “extrapolation in time” scheme as the one used in the computation of the departure point coordinate (Hortal, 2002 and Diamantakis and Vana, 2022) to estimate the divergence at the arrival point. Then the α coefficients are estimated at the mid-point of the SL trajectory, i.e. as the mean between the divergence of the current time step at the departure point and the extrapolated divergence at the arrival point.

This new formulation of α has been tested for different CATRINE test cases, but the impact of this modification is always found very small.

3.5 A robust linear and adjoint model for COMAD

It was mentioned at the beginning of this section, that the COMAD interpolation scheme was originally developed for the nonlinear version of the IFS (the forecast model) and not for the tangent linear model and its adjoint (TL/AD) which are used in the 4D-VAR system during minimization. Subsequently and before the start of CATRINE, a TL/AD option was developed but testing exposed stability problems. This issue has led us to believe that a re-formulation of the COMAD coefficient α (defined in sec 3.2), which smooths the abrupt change in the value of its derivatives when α is restricted in the interval (0,2) (see sec 3.2), will be required. Hence, this work was proposed in Task 1.2 of the CATRINE proposal.

In WP1 of CATRINE, alternative smoother formulations for the subroutine that defines the TL counterpart of the α coefficients were produced and tested, however, these were not necessary because eventually the source of the stability problems was found to be in a subtle computer code error (bug) related to the parallel implementation of the COMAD scheme and a further one in the adjoint code of COMAD. Apparently, the first one was introduced after a code optimization in the IFS code which was necessary to improve its computational performance on more recent super-computing platforms. The relevant parts of the code were corrected and tested to confirm its validity and robustness. Such tests included Taylor tests and adjoint tests at different resolutions which confirm correctness of the TL/AD. Furthermore, it was tested in long 4D-VAR experiments where COMAD had been activated in moist tracers. All completed successfully and were validated against a range of observations which confirmed good performance of the COMAD scheme.

This recent development has given us the additional capability of testing the COMAD scheme within the experimental inversion system OSSE reported in D1.1 of CATRINE. As explained in the D1.1 report, this inversion system uses an ensemble 4D-VAR data assimilation approach which relies on the use of the TL/AD code.

¹ The **CFL number** (Courant–Friedrichs–Lewy number) is a dimensionless value which gives the number of grid-lengths crossed by a parcel of air in one time step in a numerical model.

4 Investigating the boundary condition of the SL scheme near the surface

4.1 Background

The treatment of a rigid boundary as the Earth surface at the bottom of the atmosphere is not well posed in a SL advection scheme. With a flux-form formulation, it is straightforward to impose no flux at the bottom boundary of a model column. With the SL formulation, a special treatment is applied near the surface and at the top of the atmosphere. In case of high order interpolation, if the number of points available around the departure point is not enough, the order of the interpolation is reduced. Also, departure points estimated below the first model level above the surface (there is no model level at the surface) are moved back to the first level above the surface and there is then no need for a vertical interpolation.

Several test cases have shown that mass errors are significantly larger when the tracer is close to the surface. For an example, see D1.1 section 5.3. This led us to re-examine the SL algorithm near the surface, investigating alternatives for improving it.

A rigid surface acts as a wall which constrains the direction of the flow. Near the surface an ascending flow is necessarily associated with an horizontally convergent flow which feeds the ascent, and a subsiding flow diverges as it reaches the surface and spreads out horizontally. In Malardel et al, 2015, it was shown that, in a case of small scale convergent/ascending flow near the surface, the pointwise formulation of the SL scheme which makes the assumption that the grid-point value is representative of an entire grid box, is not able to resolve the horizontal transport of air at the lateral boundaries of a mesh. The COMAD corrections are designed to alleviate this problem, but either the formulation of the COMAD correction is underestimated or there are other, yet unknown, sources of error close to the surface that should be taken care of. This work started in WP1 of CATRINE but will be concluded in WP2.

4.2 Introducing an extra-level below the surface

As explained in the previous subsection, there is no model level at the surface and the computational domain for the SL advection ends at the lowest level above the surface. As a result, in certain meteorological situations of converging winds, the surface boundary can effectively act as an “infinite source of mass” for the levels above. In this subsection, we attempt to explain this issue in more detail, suggesting also an alternative way for applying the boundary condition for the SL advection algorithm at the surface.

Let NL be the lowest model level above the surface and η_{NL} its corresponding height in terms of the hybrid η vertical coordinate used in the IFS which varies between $\eta = 0$ (at the top of the model) and $\eta = 1$ (at the surface). If the vertical coordinate of the departure point η_D is such that $\eta_D > \eta_{NL}$ then the SL algorithm resets $\eta_D = \eta_{NL}$. In this case, considering a 1-dimensional vertical flow for simplicity, as the vertical location of the departure point will coincide with the location of the vertical level η_{NL} no vertical interpolation takes place, the tracer concentration after advection will be $\phi_D = \phi_{NL}$ where ϕ_{NL} the tracer value at NL. The same happens for a 3-dimensional flow except that ϕ_{NL} has been interpolated in the horizontal location of the departure point at NL to consider horizontal advection.

Therefore, given the above analysis, when the meteorological situation persistently results in departure points crossing below NL, such as in the bubble cases examined in section 3, then an artificial growth of mass occurs. The tracer value at NL remains unchanged and gradually spreads upwards. We saw that the COMAD interpolation weakens this growth, but it cannot stop it, especially the standard 2-point formulation used in cubic COMAD interpolation. Modifying the boundary condition is a possible way to tackle this. To do this we introduce an extra level NL+1 which can be symmetric to the rigid surface boundary or can be exactly at the surface boundary. $\phi_{NL+1} = 0$ is forced there and the departure point is allowed to enter the zone between NL and the surface. Linear interpolation is applied in this layer, as it is the

CATRINE

standard practice near the boundaries, due to lack of sufficient amount of data to construct a higher-order interpolator. Therefore, the situation $\phi_D = \phi_{NL}$ no longer occurs but instead $\phi_D = \gamma\phi_{NL}$, where γ is the interpolation weight defined as

$$\gamma = \frac{\eta_{NL+1} - \eta_D}{\eta_{NL+1} - \eta_{NL}} < 1.$$

This effectively means that the infinite supply of mass from the boundary has been stopped. The outcome of such a modification can be seen in figure 4.2 where a bubble case like the one in sec 3.4.1 is shown. The bubble is defined at the IFS cubic grid (see Malardel et al, 2015) at 125 km grid-spacing, with 137 vertical levels and the simulation runs with 1h timestep. The resolution differs with the corresponding resolution of the case in sec 3.4.1, however, qualitatively the bubble evolves the same way. The plot shows the result of the simulation after 96 timesteps. When the standard SL advection is used, there is fast growth of mass while the value at the bottom remains undiluted i.e. the bubble is lifted but leaves a trail of mass below. This unphysical effect weakens considerably when applying COMAD interpolation: the mass growth is slower and while the bubble again does not detach from the surface it looks more physical. When the alternative boundary condition is applied to COMAD SL interpolation the mass growth is even smaller, and the bubble fully detaches from the surface.

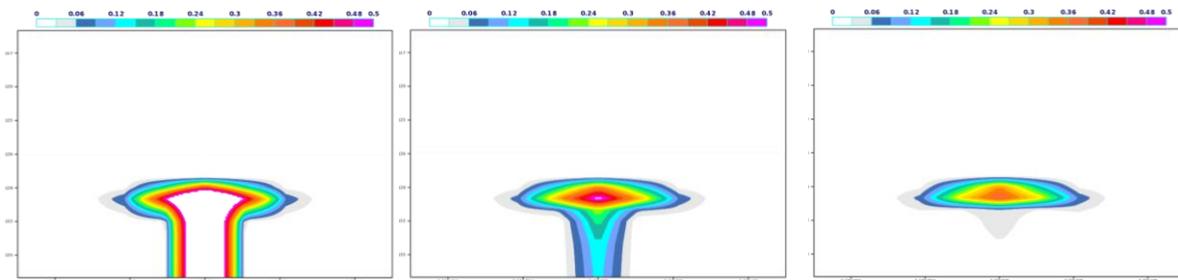


Figure 4.2. Bubble tracer case simulation. A zonal vertical cross-section slice of the tracer after 96 time steps. Left: standard SL advection with cubic interpolation and LQM3D limiter. Middle: SL with COMAD and LQM3D limiter. Right: as in the middle but assuming existence of an extra level below the surface with 0 tracer concentration. The white colour indicates that the tracer concentration is above the contour maximum 0.5. The initial bubble tracer concentration is 1. The vertical axis shows the IFS model level number (1-137) and the horizontal the longitude (from 4° W to 4° E).

Although this method gives very good results in “bubble” cases it cannot always provide satisfactory results in general cases as finite-volume schemes do; essentially, although it achieves a similar outcome for bubble cases, it is not fully consistent with the three-dimensional finite-volume / donor cell analogy - it tackles it through a vertical fix which works only in the case of an ascending flow with “polluted” air. It treats a very specific weakness of the SL advection schemes which occurs at specific meteorological situations when a mass source is near the boundary or emitted from the surface. Further testing of this alternative algorithm (see sec 5.3) shows overall reduction of mass conservation errors. However, when the tracer fills the entire atmosphere, it tends to reduce its overall mass, rather than increase it which is the typical behaviour with the standard SL algorithm, a drawback which seems to be partly regulated by the limiter. The question to address during the early stages of WP2 through additional experimentation and evaluation is if this is merely a method for compensating positive mass error or can it have a genuine positive impact in tracer transport across different weather regimes.

4.3 Modifying COMAD weights near the surface

The extra level below the surface, can be applied either in standard Lagrange or COMAD interpolation. Other alternatives for improving COMAD interpolation near the surface have also been investigated which we summarize in this short section. The motivation was the assumption that COMAD underestimates the correction because of the simplicity with which we represent the deformation of a material element in the fluid. In particular, the COMAD

CATRINE

correction is based on a symmetric deformation of the distance element around its midpoint along the direction of interpolation during a time step. Close to the surface, the deformation of the flow is strongly constrained by the presence of a rigid boundary. The hypothesis of symmetry around the midpoint no longer holds as the flow deformation is mostly driven by the transition from horizontal to vertical motion or vice-versa. Several attempts have been made to modify the COMAD weights to better account for the real flow deformation near the surface.

One of them was to modify the vertical divergence used in the COMAD vertical correction, assuming that the vertical velocity is zero at the first full level rather than at the surface (which is a half-level). For consistency, correcting the vertical divergence used in the vertical COMAD weights also requires correcting the horizontal divergence, so that the sum of the horizontal and the vertical divergences remains unchanged. Therefore, corrections are applied to both horizontal and vertical COMAD weights. This solution works well for a tracer initially placed only in the first layer above the surface and for linear interpolation. With cubic interpolation, there is a systematic mass loss. We will continue to investigate possible changes in COMAD and we will conclude on their usefulness during the coming months at the beginning of WP2.

5 Limiters with improved conservation properties for tracers

The advection process should not generate new minimum/maximum values when transporting a field. However, a SL scheme which uses an interpolation method with order greater than 1, can generate new minimum/maximum values (overshoots/undershoots respectively) and even negative values when the advected field has strong gradients. This is often the case in tracers which are relevant to the CATRINE project such as greenhouse plumes emitted from anthropogenic sources as well as other tracers such as cloud fields, aerosols and volcanic plumes. The IFS advection scheme uses interpolation limiters (also known as monotone filters) to prevent this problem enforcing monotonicity in advection and ensuring that concentration values for tracers remain always positive. It avoids unphysical oscillations in the solution with a small penalty in accuracy as it is only activated locally where needed without enforcing strict monotonicity in the entire domain of the interpolation polynomial function. Such schemes are often quoted as “quasi-monotone” SL schemes.

The IFS is using a 32-point interpolation stencil to perform a 3D-cubic interpolation at a departure point. This is simplified to a set of separate 1D-cubic interpolations within the stencil which are performed in the zonal, meridional and vertical directions (Ritchie et al, 1995). The limiter used by the IFS advection scheme is based on the work of Bermejo and Staniforth (1992) where the three-dimensional quasi-monotone limiter (QMSL) is introduced. There are two versions of this limiter in the IFS code: the “classic” QMSL limiter which is three-dimensional, and it is applied at the end of all cubic interpolations performed in the 3D interpolation stencil i.e.

$$\phi^*_D = \max(\phi_{min}, \min(\phi_{max}, \phi_D))$$

where ϕ_{min} , ϕ_{max} the local minimum or maximum values of the 3D box which is formed from the 8 points that surround the departure point D, ϕ_D is the result of the interpolation and ϕ^*_D the result after the limiter has been applied. This limiter will be labelled as “LQM3D” and its function is to ensure that the final (limited) interpolated value ϕ^*_D lies between the local minimum and maximum. A split-dimensional version of this limiter, known as “LQM”, is additionally available in IFS. This also clips to the minimum / maximum; however, its difference is that it is applied at the end of each 1D high order interpolation within the stencil. To build a cubic interpolant, the IFS does 7 cubic interpolations (see Ritchie et al, 1995) and therefore, this limiter is activated multiple times resulting in stronger damping of the advected field which could be beneficial for conservation but can also produce more diffusive results.

In D1.1 we explained that the limiter modifies and destroys the conservative properties of the Lagrange polynomials that the SL advection relies on. We also demonstrated that the limiter introduces significant mass conservation errors when localised tracers with sharp gradients such as plumes are considered. Hence, a clear recommendation was to work towards an improved, more conservative but at least equally accurate version of the existing limiters.

Our first attempt was to investigate the ILMC (Iterative Locally Mass Conserving) monotone filter for SL advection (Sorensen et al, 2013). This is a three-dimensional conservative adjustment of the LQM3D limiter which ensures that the global mass of a tracer after applying the ILMC limiter (filter) is equal to the mass after the interpolation but before the limiter is applied. It is achieved by a local adjustment of tracer concentration values, in grid-points where overshoots or undershoots have occurred: the mass clipped is given to or taken from neighbouring grid-points if these have the “capacity” to accept such change without violating their own local monotonicity criteria. However, an implementation of this work in IFS, suitable for parallel computers, would have resulted in non-reproducible results when changing the number of processors that a simulation is run. Furthermore, there were technical difficulties in implementing this scheme in a mature operational system such as the IFS which uses specific data structures to achieve a very efficient, second level of parallelization with openMP across different high-performance computing (HPC) architectures. In the existing IFS software design,

CATRINE

the direct memory access to grid-point neighbours, required by the ILMC algorithm, is not straightforward and it has an additional detrimental impact on computational efficiency which is undesirable given that this algorithm is by design already computationally expensive.

To circumvent the implementation problems described in the above paragraph, modified versions of the ILMC filter which adjust the mass in a vertical column locally rather than in a three-dimensional manner were considered. This approach was also motivated by the fact established in D1.1 that mass conservation violations in the vertical are the most significant and therefore a vertical adjustment is a reasonable approach. Two approaches were considered: the first applies to LQM3D limiter and it will be labelled in what follows as “LQM3DCONS” while the second applied to the LQM limiter and it will be labelled as “LQMVCONS”. Both are simple algorithms which are activated once a violation of local monotonicity is detected. As in ILMC, the amount of mass clipped is transferred in the neighbouring points but the adjustment is restricted within a column only. The difference between the two new proposed approaches is that LQMVCONS distributes only the clipped mass by the vertical part of the limiter while LQM3DCONS distributes the entire clipped mass. Experimentation with applying the correction to the neighbouring column point sweepings upwards or downwards was done which showed that the latter was more accurate and more conserving. We were able to quickly discard the less suitable option, before we conduct expensive long CATRINE simulations, using the NWP system and experimenting with moist tracers which allows us to assess model changes verifying against ECMWF analyses. This is another strength of the IFS system. Finally, as we will show in the experiments of section 5.2 and 5.3, we have also applied these new limiters iteratively: first in a downward direction (starting from the top) and then repeating the mass adjustment in an upward direction.

5.1 Impact of limiters in tracer correlation test case

To assess the impact on accuracy and the conservation properties of the new limiters the DCMIP12 test case 11 on correlated tracers (Kent et al, 2014) was introduced in the IFS. In this test case, the 2D correlated case proposed by Thuburn and Lauritzen (2012) is expanded in a more challenging three-dimensional deformational flow setting. It tests the ability of advection schemes to preserve non-linear correlation relationships between tracers and can identify situations in which unphysical numerical mixing occurs. These are important properties for atmospheric composition forecasts such as the ones considered in CATRINE. It also gives useful information on accuracy aspects of tracer transport, for example it exposes schemes that are over-diffusive.

CATRINE

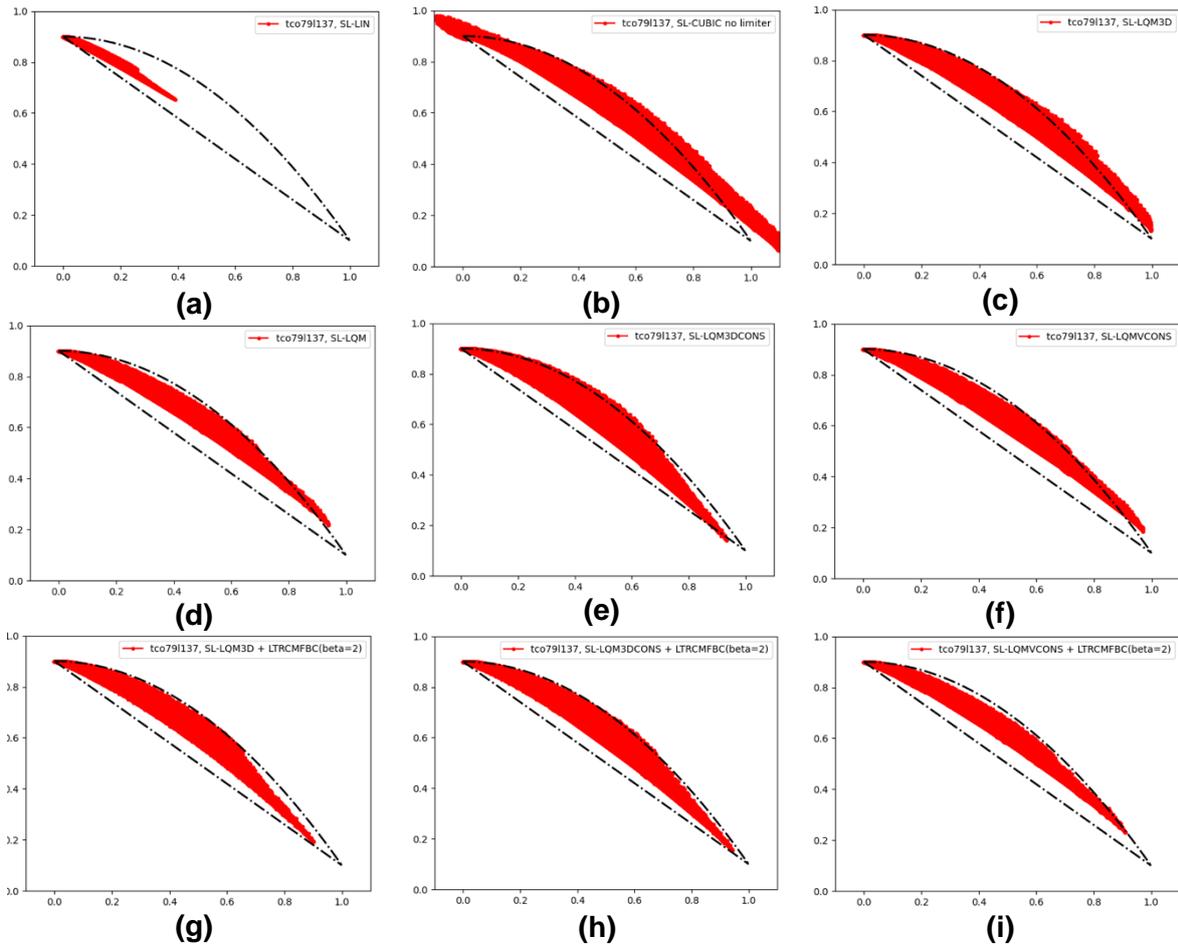


Figure 5.1. Tracer correlations for different options of the SL advection scheme. Option (a): Linear interpolation is used; (b) cubic without limiter (c): cubic with LQM3D limiter; (d): cubic with LQM limiter; (e): cubic with LQM3DCONS limiter; (f): cubic with LQMVCONS limiter; (g,h,i): cubic with LQM3D, LQM3DCONS and LQMVCONS limiter respectively and mass fixer configuration used in WP7 tests of D1.1.

The performance of the IFS advection scheme on this test is illustrated in the scatter plot of figure 5.1. The grid used for this test is tco79 (approximately 125 km grid-spacing) with 137 vertical levels which is as close as possible to the recommended resolution in the paper. A desirable result in these scatter plots, where the concentration of one tracer is plotted against the concentration of the other tracer, is to have the “red area” confined within the dashed-dotted enclosed shape following the parabolic curve. Red areas outside this region suggest that unphysical numerical mixing can occur and when the red points are gathered mostly at the top part of the enclosed conic shape excessive diffusion occurs. The SL advection scheme based on linear interpolation exhibits such excessive diffusion (figure 5.1a) while when cubic interpolation without limiter is used strong unphysical numerical mixing occurs (figure 5.1b). Use of a limiter improves results by reducing the numerical mixing. The LQM limiter (figure 5.1d) seems to produce better results than the LQM3D (figure 5.1c) but LQM3DCONS is slightly better than LQMVCONS (compare figure 5.1e versus figure 5.1f). Furthermore, use of the mass fixer seems to have an additional beneficial impact (figures 5.1g,h,i) and the mass fixer combined with LQM3DCONS limiter (figure 5.1h) yields the best results (no numerical mixing and little diffusion). Therefore, the new limiters improve tracer results especially when combined with the mass fixer resulting in a scheme which exhibits a small amount of numerical diffusion and preserves reasonably well correlations between tracers.

5.2 Plume tracer passive advection case study

The new limiters have also been tested and compared against the standard ones in IFS (LQM, LQM3D) on the test cases considered in D1.1 (sections 5.1 and 5.2). As explained in that deliverable, these cases are very challenging tests for the conservation properties of the SL advection scheme, because of the very sharp gradients the tracers exhibit, the lack of background values of tracer and the fact that no further physical process such as turbulent diffusion acts on them which can smooth tracer gradients. The expectation is that these new limiters will reduce the conservation error, therefore the mass fixer will compute a smaller correction.

In figure 5.2, results from the plume tracer case described in section 5.1 of D1.1 are shown. In this case, which is based on the work by Eastham and Jacob (2017), a rectangular tracer with dimensions 4x5 degrees (longitude x latitude), is initialized in locations in the mid-latitude and the tropics. In the interior of the rectangular tracer, the concentration (mixing ratio) increases abruptly to a fixed value while outside is 0. The plume tracer is passively transported by the winds predicted by the IFS model, which is initialized from the ERA5 analysis of 01/07/2022. We can notice there that the limiters have a large impact on the mass conservation error growth which is much higher when the tracer is initialised at the surface level. This point was analysed and discussed in D1.1. The LQM limiter (blue line) produces smaller conservation error than LQM3D (black line) and both new two limiters substantially reduce the mass conservation error. The LQM3DCONS (red line) is the most conserving for the tracer initialised at the surface, for all three locations. It reduces the conservation error by a factor that ranges between 4 and 7. It also shows a large reduction of conservation error for the tracers initialised at the upper boundary layer but less for the upper stratosphere case which is the one where the smallest errors occur. LQMVCONS performs also well, it shows a smaller reduction compared to the basic scheme LQM which has been derived from, simply because the latter is already more conserving. Adding an extra iteration (dashed lines) in these algorithms reduces further the conservation error and, in some cases, it nearly eliminates it. However, the extra iteration may result in mass spreading in the column further away from the point that the overshoot/undershoot occurred, which could compromise the local nature of the algorithm. This is more noticeable for tracers initialised at levels far away from the surface.

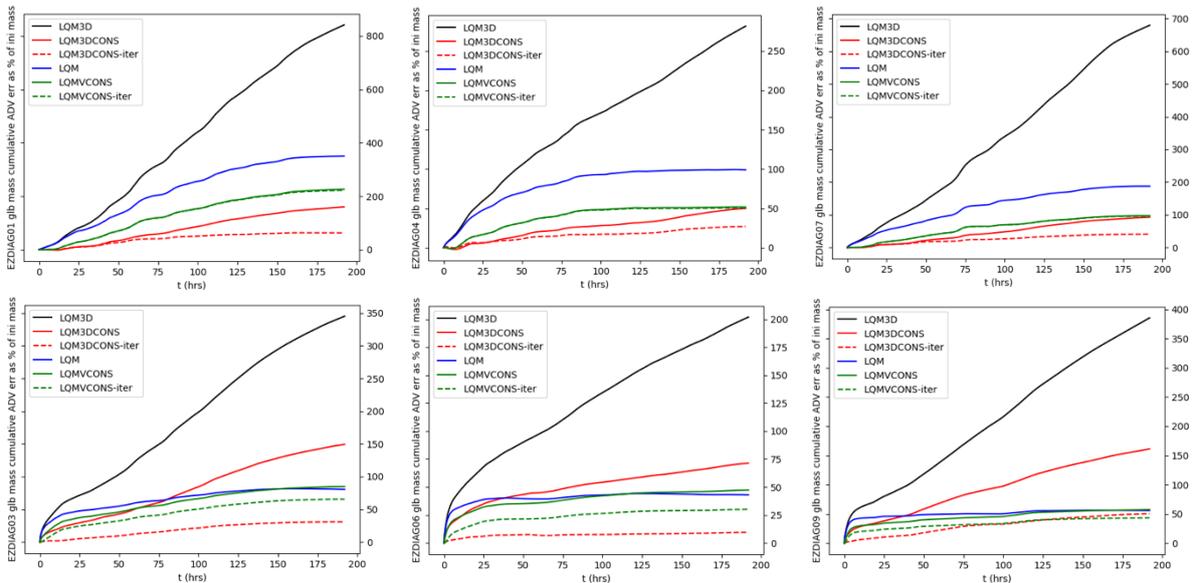


Figure 5.2. Mass conservation error as a percent of the initial global tracer mass, for the passive advection “plume tracer” case study. Top row: tracer initialized at the surface level. Bottom row: tracer initialised at approximately at 900 hPa (930m height). Left column: location in eastern China. Middle column: Location in Indonesia. Right column: location in the east coast of the USA. COMAD with four cubic interpolation limiters is tested: LQM, LQM3D, LQM3DCONS, LQMVCONS. The last two apply the limiter using an extra iteration and have been labelled as LQM3DCONS-iter and LQMVCONS-iter.

5.3 Testing the limiters on tracers emitted by point sources

The point source tests described in section 5.2 of the CATRINE D1.1 report have been applied here to assess the newly developed limiters. In these tests, different tracers are emitted from point sources at various geographical locations. Each emission increases the concentration of the corresponding tracer by 0.5 kg/kg per timestep at the grid-point representing the source, which lies at the model level closest to the surface. The simulations use the Tco399 model grid (with an average grid spacing of 28 km) but our results remain qualitatively consistent when higher-resolution grids are employed. The emitted tracers are transported through the atmosphere by forecast winds from the IFS experiment. The forecast is initialized from ERA5 analysis on 01/07/2022; winter days have also been considered but they are not shown here as they lead to similar conclusions. As in the plume case of the previous sub-section, advection is passive so the only transport process that acts on these tracers is 3D-advection without parametrized turbulent diffusion and convection. This helps us to study purely the impact of the dynamics in the growth of mass conservation errors.

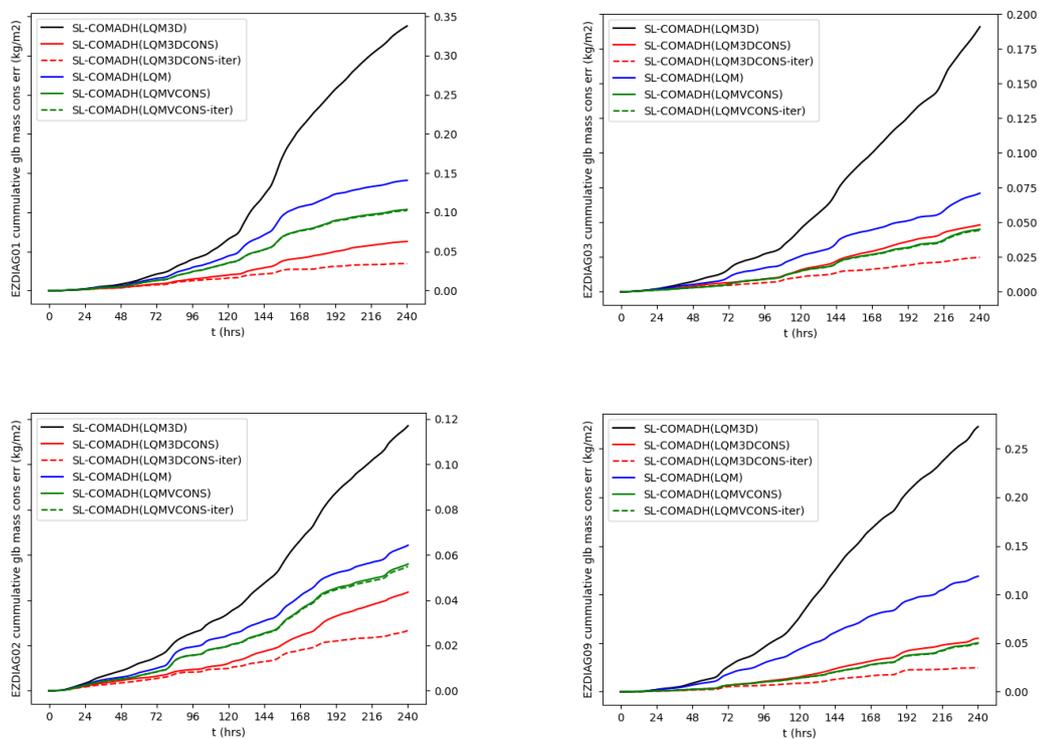


Figure 5.3a. Accumulation of mass conservation error in kg/m² for a case where a tracer is emitted from a surface point at different locations near industrial zones and advected by the atmospheric flow. Top left: point source in China with coordinates (113E,37.5N); Top right: point source in the USA with coordinates (77W, 40.5N); Bottom left: point source in Indonesia with coordinate (113E,7.5S); Bottom right: point source in South Africa (115.4E,20.8S). The same limiters with those in figure 5.2 are tested applied to the latest improved version of the horizontal COMAD interpolation scheme.

The results obtained confirm the significantly better conservation properties of the newly developed limiters. This is illustrated in figure 5.3a where the accumulated mass conservation error in kg/m² is plotted for four tracers emitted each at four different locations. The results show that the LQM3DCONS limiter, compared with LQM3D, reduces mass conservation error by a factor between 3 and 5 while LQMVCONS, compared with LQM, reduces the mass conservation error by a noticeably smaller factor ≤ 2 . This is because LQM produces much smaller conservation error and LQMVCONS adjusts only the conservation error produced by the vertical part of the limiter. Again, as in the previous case, if an extra iteration is added the

CATRINE

mass conservation error reduces further in LQM3DCONS while the impact on LQMVCONS is negligible.

Finally, we have also used this test case to experiment with the alternative boundary condition described in sec 4.2. Indeed, as seen in figure 5.3b, it reduces the magnitude of the growth in mass per unit area caused by conservation errors by 35% to 40% depending on location. It is worth mentioning, that if the point source is placed at a higher level than the lowest atmospheric level, as in sec 5.3 of CATRINE D1.1, there is very little or no impact at all from this alternative boundary condition; as expected the boundary condition impact diminishes as the tracer is moved away from the surface.

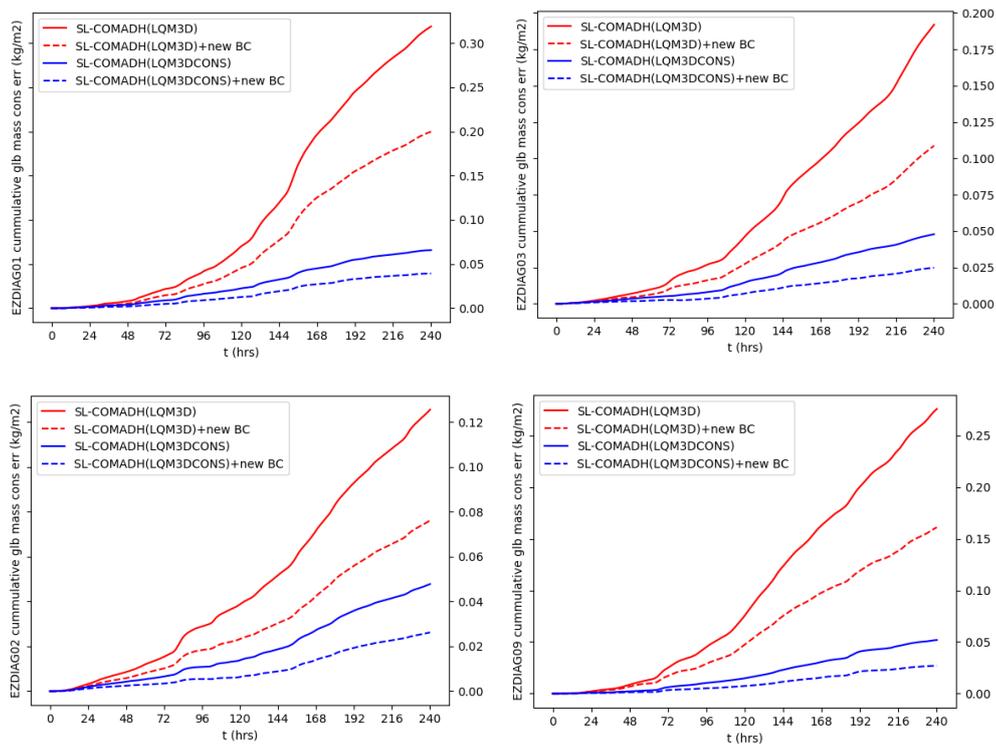


Figure 5.3b. Accumulation of mass conservation error in kg/m² for the same case as in figure 5.3a for LQM3D and LQM3DCONS with standard and alternative boundary conditions (labelled as “new BC”) as described in sec 4.2.

6 Advection of density versus specific ratio

6.1 Background

One of the main reasons for the lack of mass conservation in the IFS is that the equations are not written in conservative flux form (see D1.1 for more detail).

Indeed, the semi-Lagrangian (SL) scheme used in the IFS requires equations to be expressed in advective form. Currently, these equations are formulated in terms of specific ratios—namely, the ratio of tracer density to the density of moist air. As a result, the conservation of tracer mass relies on two components: the continuity equation for total air mass (Eqn 1, where m is a pseudo-density, a quantity equivalent to density in a hybrid vertical coordinate framework) and the equation for the specific ratio q with respect to total mass (Eqn 2). Any mass conservation error can thus originate from inaccuracies in either equation.

$$\frac{Dm}{Dt} = -m (\nabla_{\eta} \cdot \vec{v} + \frac{\partial \dot{\eta}}{\partial \eta}) \quad (1)$$

$$\frac{Dq}{Dt} = 0 \quad (2)$$

$$\frac{D(mq)}{Dt} = -(mq) (\nabla_{\eta} \cdot \vec{v} + \frac{\partial \dot{\eta}}{\partial \eta}) \quad (3)$$

To assess the impact of this separation on mass conservation errors in SL tracer advection, we modified the code to instead use an advective-form equation for the tracer density directly (Eqn 3). Thus, the transport, which is a non-linear operation, is directly computed for the mass of the tracer, not anymore for the ratio with respect to a reference mass. Unlike the specific ratio formulation, this new tracer density equation includes a right-hand side (RHS) term—similar to that of the total mass equation—which controls the fastest waves in the hydrostatic system (gravity and Lamb waves). In the case of the continuity equation, this RHS is stabilized by the semi-implicit (SI) solver.

To ensure consistency in mass continuity treatment, stability and maintain computational efficiency, the RHS of Eqn 3 is computed *after* the SI solver step.

Another reason for this modification is that the COMAD interpolation weights (Malardel and Ricard, 2015) theoretically improve conservation for equations written for a variable per unit volume (density equivalent). In practice, they also improve conservation when specific ratios are used, despite exhibiting asymptotic behaviour at stationary, no-wind points, which causes them to lose their conservative properties.

6.2 Results

Equations (3) have been coded in the IFS for tracers. They have been validated with a selection of the stress test cases described in detail in D1.1.

Figure 6.2 shows the mass conservation error in the bubble case (left panel) and in the mountain case (right panel) with tracer near the surface for both the advection of the specific content of tracer (equation without RHS) or the mass or tracer per unit volume (equation with a RHS). For both cases and for both linear and cubic interpolation, the mass conservation error is independent of the variable which is advected.

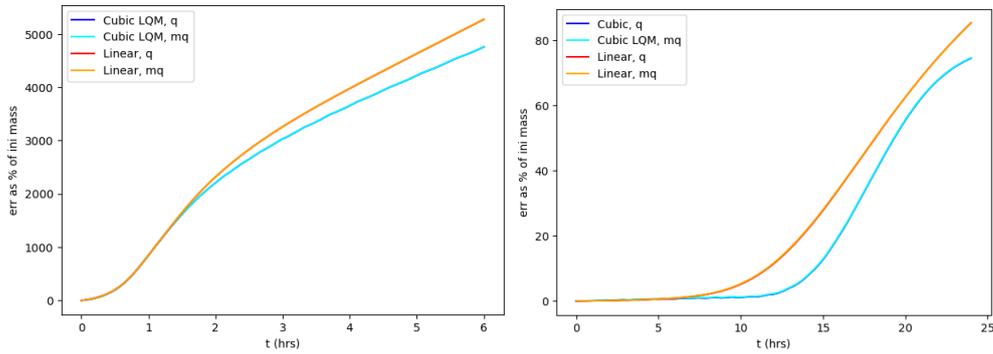


Figure 6.2: Mass conservation error in the bubble case (left) and the mountain case (right) for cubic and linear advection in the SL scheme. Curves for the advection of the specific ratio of tracer (q) or the advection of the mass of tracer per unit volume (mq) are superimposed.

These results demonstrate that using ratios rather than density in the IFS is a negligible source of error compared to other sources. The lack of conservation is then primarily due to the advective form of the equations, for which it is not possible to exert exact control over the input/output of mass in a grid box.

7 Testing the improvements using WP5 and WP7 protocols

Global model emission protocols for model inter-comparison studies have been defined in CATRINE WP5 and WP7. We have used these protocols to assess the performance of our most promising new developments in terms of conservation error but also in terms of the overall transport accuracy. For the latter we have used the optimized WP7 CO₂ fluxes. The experiments performed here, which use WP7 protocols, use the same setup as the experiments of D1.1 with the difference that the new limiters and the improved version of COMAD described in section 3.2 has been used. Results from these tests and the previous case studies as well as the knowledge we have built from various analyses presented here and in D1.1 have helped us to make recommendations for the most beneficial modifications for the global CO2MVS model IFS which will be finalised in WP2.

In figure 7a results from a simulation that follows the WP5 protocol based on CAMS fluxes is shown. This setup has been used since the early stages of this project for assessing the sensitivity of different components of the advection scheme, such as limiters, with respect to mass conservation. An example of the total column mass conservation error of CO₂ (from the advection scheme) for a period of approximately 6 months is shown in figure 7a. This is computed by taking the difference of a simulation without mass fixer, in which the mass conservation error is let freely to grow, from an equivalent one which uses mass fixer to keep the global tracer mass constant during SL advection. In the left plot of figure 7a, we compute this for a simulation which applies LQM3D limiters and the COMAD scheme in the horizontal. In the right plot of figure 7a, a simulation using LQM3DCONS and the enhanced COMAD (applied in both convergent and divergent flows as in sec 3.2) is used. The reduction of the mass conservation error here is nearly five times. We have repeated similar simulations applying LQM and LQMVCONS, however, it turns out that for the real CO₂ tracer, which is well mixed with the background air, the LQM3D and LQM3DCONS limiters produce better results. In this situation these two limiters can be more accurate as they are less active, and the amount and strength of monotonicity violations are not as extreme as the single plume cases previously examined.

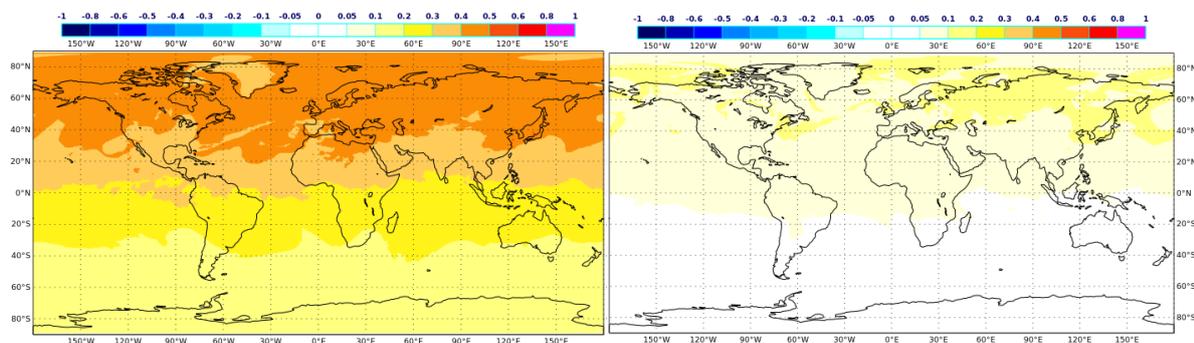


Figure 7a. Total column mean molar fraction change (in ppm units) of CO₂ field due to SL advection mass conservation error in a WP5 simulation that runs between 20/12/2020 until 1/06/2021. Left: with standard LQM3D limiter. Right: LQM3DCONS and improved COMAD. The mass growth is deduced comparing a run without mass fixer against an equivalent run with mass fixer. The global mean and RMS of the total column change in the left plot are 0.296, 0.313 while for the right 0.061, 0.066.

In section 5 we demonstrated that the newly introduced limiters reduce significantly the mass conservation error that the IFS advection scheme produces. However, in section 6 of D1.1, it was also established that in real greenhouse gas simulations, such as those following WP5 and WP7 protocols, the mass conservation error is several orders of magnitudes smaller than

the one produced by our stand-alone plume tracer cases of section 5 of D1.1 and of this report. It was explained that the main reason for this behaviour is much weaker tracer gradients in real simulations as the greenhouse tracers are well mixed with the background air, and because of weaker emissions than those considered in our artificial plume tests which have been designed to stress tests our algorithms and to help us diagnose conservation issues. Furthermore, in real simulations, we use the mass fixer which inhibits growth of mass and the growth of tracer gradients and additionally, transport is no longer the result of passive advection as physical parametrizations such as turbulent diffusion and convection modify tracer fields weakening further these gradients. A question worth exploring here is if there is evidence of reduced conservation error in the WP7 framework from these recent model improvements and what is the sensitivity of the model to such conservation improvements with respect to its ability to reproduce observations.

The first question is explored in Figs. 7b, 7c. Time series of the tracer mass conservation error for WP7 simulations with different limiters is shown. This error is measured by a mass fixer diagnostic, and it is equal to the total correction it computes to restore conservation. The plots in figure 7b concern the anthropogenic CO₂ tracer field (CO2APF) emitted within the forecast period of our simulations which starts on 01/12/2021 and ends on 31/12/2022. CO2APF is only a component of the real CO₂ field. It starts with zero concentration at the beginning of the simulation, but its concentration gradually increases and mixes with the remaining background CO₂ and the air. The left plot of figure 7b, compares experiments with different limiters, for a 24-hour period starting on 02-12-2021, near the start of the simulation. The lack of mixing at this early stage of the simulation implies sharper gradients which is reflected in the larger mass conservation errors seen, compared to later dates (as those in the right-hand plot of Fig 7b). The conservation error for CO2APF is smaller with LQM3DCONS while all the other limiters produce roughly the same conservation error. This limiter is also compared against LQM3D and LQM for later dates, approximately 6 and 12 months after (see right plot of Fig 7b) where we notice that it is also the best in terms of conservation error, reducing it by a factor between 2-3. For conciseness in the plot, the corresponding time series from the LQMVCONS limiter have not been included here as these showed no advantage in this type of experiment. As seen in the D1.1 WP7 simulations, the LQM and LQM3D limiters have similar conservation errors and often LQM3D seems better than LQM. In figure 7c the same comparisons are shown for the CO₂ and CH₄ tracers. Again, simulations with LQM3DCONS have the smallest conservation errors, the conservation error reduces approximately by a factor 2-3.

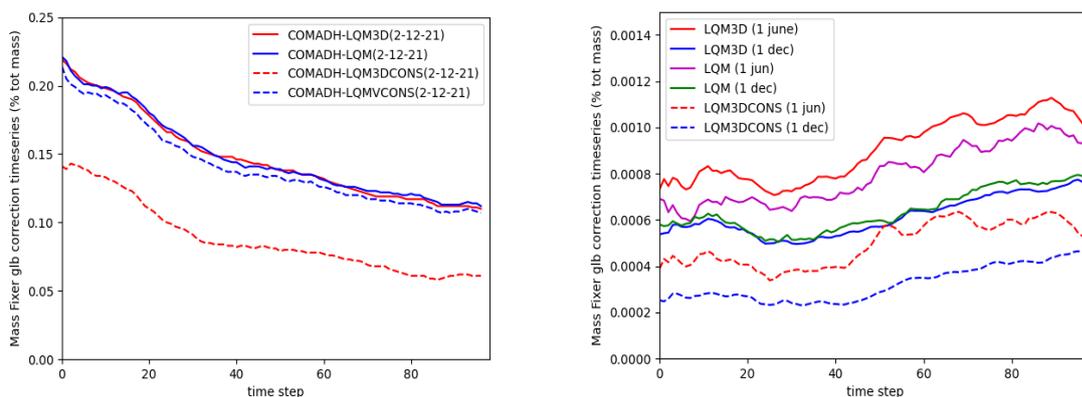


Figure 7b. Time series of mass fixer correction (in percent of the total tracer mass) for the CO2APF tracer (anthropogenic CO₂) at each time-step from WP7 simulations when different limiters are used. Left: Time series from a date at the beginning of the simulation (02-12-21). Right: Time series for two later dates (1 June 2022 and 1 December 2022) for different limiters.

CATRINE

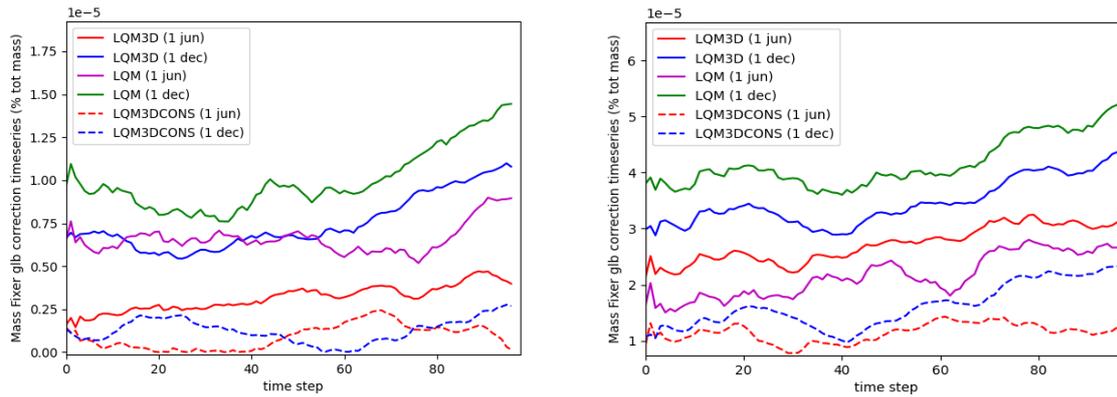


Figure 7c. Time series of mass fixer correction (in percent of the total tracer mass) at each time-step from WP7 simulations when different limiters are used. Left: CO₂. Right: CH₄.

In figure 7d the daily average XCO₂ (column averaged CO₂ mixing ratio) model time series from IFS experiments are plotted against TCCON observations for four different locations. The experiments compared are: (i) an experiment which uses the LQM3D limiter and the mass fixer² with $\beta = 2$ (light blue line); (ii) an experiment with the same limiter but with mass fixer with $\beta = 3$ (red line); (iii) an experiment with the LQM3DCONS limiter without mass fixer (purple line); and (iv) an experiment with LQM3DCONS limiter and the mass fixer with $\beta = 2$ (blue line). At the top of each plot the bias (δ), standard deviation of errors (σ) and model-observation correlation coefficient (r) error metrics (against observations) are indicated. These are averaged in the period of simulation (13 months). As these plots show, there is good agreement between the model and the observations which are indicated by the thick black dots. When the mass fixer is used, there is little sensitivity with respect to limiters which agrees with results from similar experiments conducted in D1.1. The LQM3DCONS limiter has a neutral impact on accuracy but as we saw from the corresponding time series plots of figure 7b and 7c, it reduces the conservation error. However, when the mass fixer is not activated, the LQM3DCONS simulation has noticeably smaller bias from equivalent simulations with different limiters (LQM, LQM3D). This is not evident from the plots in this report but it can be verified by comparing the biases in Table 1. This table summarizes results from the simulations in Figure 7d and from additional simulations which took place earlier in WP1 and were documented in D1.1. So, all our results so far suggest that LQM3DCONS reduces the conservation error with neutral or even positive impact on accuracy (as seen by the experiment without mass fixer).

² As explained in D1.1 the mass fixer β parameter determines the magnitude of the mass fixer correction weight which is proportional to the tracer gradient, larger β values increase the magnitude of mass fixer correction in regions of steeper gradients while they reduce it to almost 0 where the tracer field is very homogenous (for details see D1.1, sec 6).

CATRINE

XCO2 - TCCON			XCO2 - TCCON		
LQM3D + MF(OFF)			LQM + MF(OFF)		
Station	δ	σ	Station	δ	σ
Karlsruhe	1.31	0.80	Karlsruhe	1.67	0.88
Xianghe	2.24	1.85	Xianghe	2.95	1.97
Izana	0.66	0.41	Izana	0.97	0.38
Lauder	0.50	0.38	Lauder	0.56	0.40
LQM3D + MF(ON, $\beta = 1.5$)			LQM + MF(ON, $\beta = 1.5$)		
Station	δ	σ	Station	δ	σ
Karlsruhe	0.59	0.79	Karlsruhe	0.70	0.76
Xianghe	1.50	1.42	Xianghe	1.73	1.89
Izana	0.01	0.46	Izana	0.03	0.43
Lauder	0.16	0.40	Lauder	0.11	0.42
LQM3D + MF(ON, $\beta = 2$)			LQM3DCONS + MF(OFF)		
Station	δ	σ	Station	δ	σ
Karlsruhe	0.55	0.79	Karlsruhe	0.75	0.81
Xianghe	1.22	1.71	Xianghe	1.42	1.73
Izana	-0.01	0.46	Izana	0.15	0.41
Lauder	0.20	0.40	Lauder	0.25	0.39
LQM3D + MF(ON, $\beta = 3$)			LQM3DCONS + MF(ON, $\beta = 2$)		
Station	δ	σ	Station	δ	σ
Karlsruhe	0.55	0.80	Karlsruhe	0.58	0.81
Xianghe	0.85	1.46	Xianghe	1.19	1.71
Izana	-0.02	0.44	Izana	0.00	0.43
Lauder	0.21	0.40	Lauder	0.18	0.40

Table 1. IFS verification statistics from CO₂ simulations with different limiters and different options for the mass fixer. The bias (δ) and the standard deviation (σ) against TCCON at four different sites corresponding to simulations in Figure 7d and additional simulations that correspond in Figure 20 of CATRINE deliverable report D1.1 are gathered here. The label MF(OFF) implies that the mass fixer for CO₂ was switched off. Units in ppm.

CATRINE

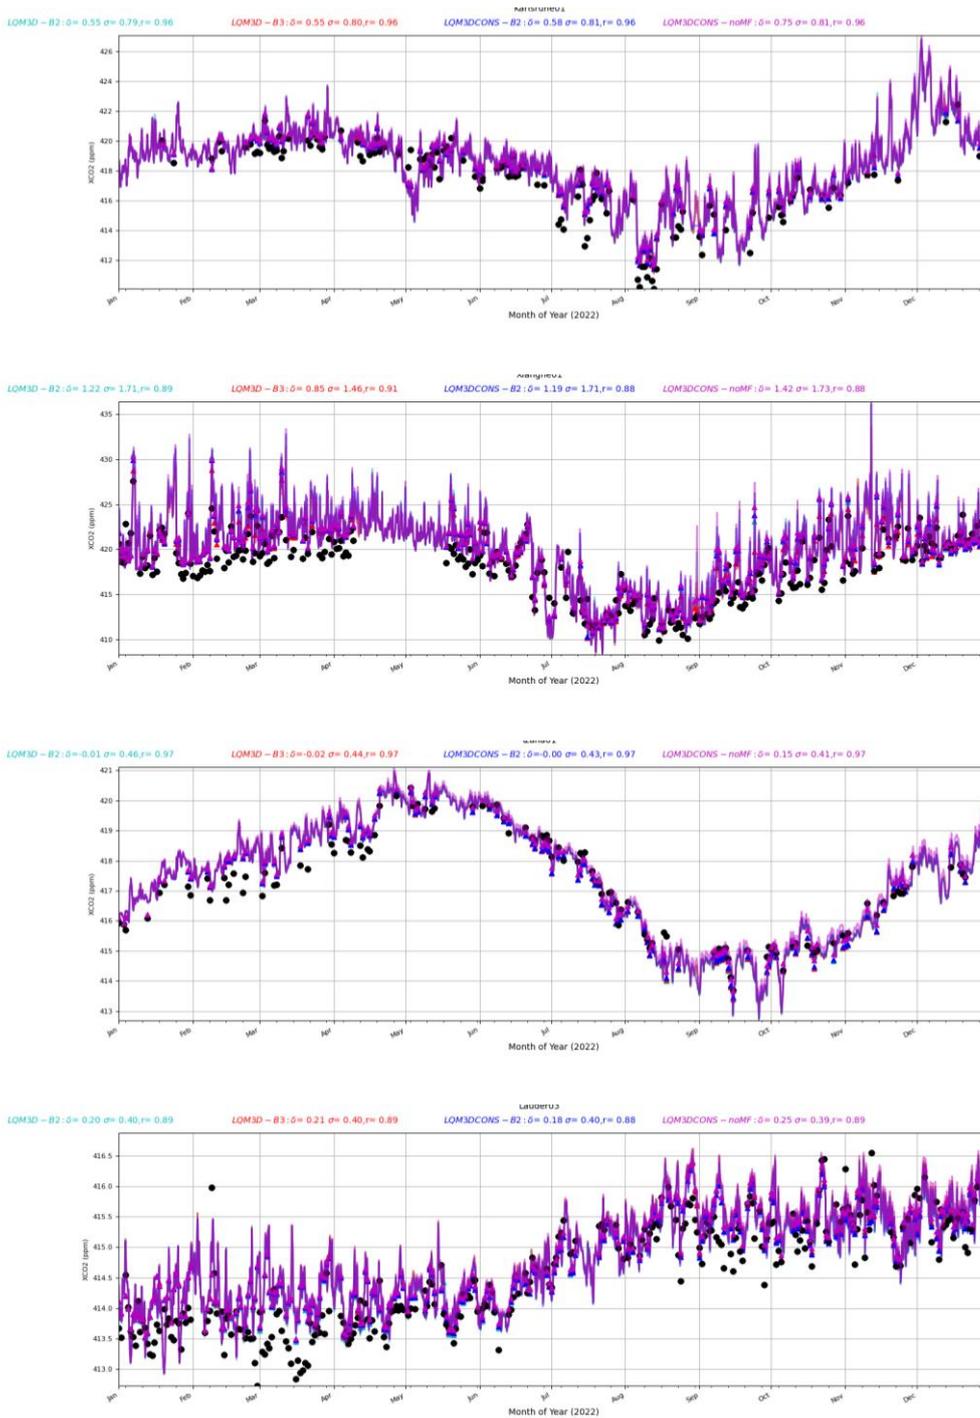


Figure 7d. Timeseries (January-December 2022) of daily average XCO₂ compared to TCCON observations at four sites from simulations with different combinations of limiters and the mass fixer or without the mass fixer (purple line labelled noMF). From the top to the bottom: near the city of Karlsruhe (Germany) in a forested area, Xiangei (China), Izana (Canary Islands), Lauder (New Zealand). The first two are near areas of industrial/anthropogenic emissions while the last two are away from pollution sites. In each plot the bias, standard deviation of errors and correlation coefficient is written at the top for each simulation. These are computed against observations (black dots). Units in ppm.

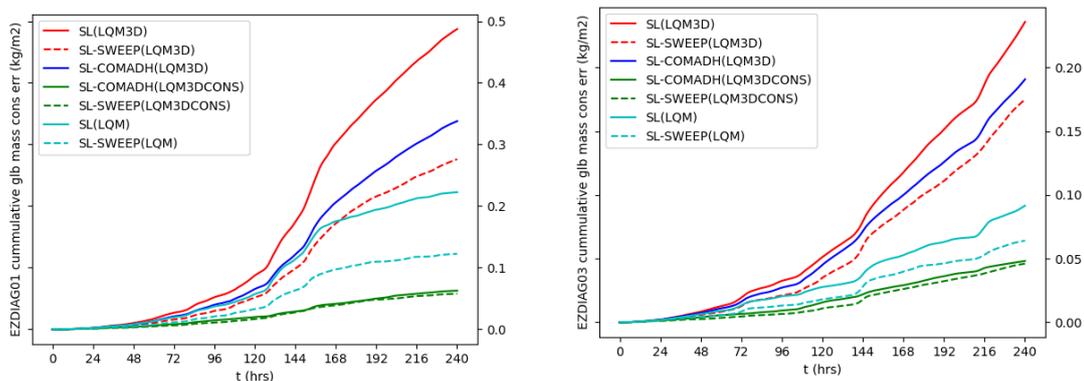
8 Sweep interpolation: a promising alternative for the semi-Lagrangian scheme for tracer transport

In their 2024 study, Mortezaazadeh et al. introduce “sweep interpolation,” a semi-Lagrangian scheme designed to enhance computational efficiency while maintaining high accuracy in numerical weather prediction models. This method alternates between quadratic, third-order backward and forward polynomial interpolations over two consecutive time steps, effectively cancelling third-order errors and achieving fourth-order accuracy with the computational cost of a third-order scheme i.e. the same accuracy with the cubic interpolation scheme we use in IFS for tracer transport at reduced cost. Implemented within Environment and Climate Change Canada’s Global Environmental Multiscale (GEM) model, sweep interpolation demonstrated up to a 15% reduction in total simulation time, particularly benefiting scenarios involving numerous passive tracers. Validation across various benchmarks, including global ozone forecasts, confirmed that sweep interpolation closely matches the accuracy of traditional cubic interpolation methods, offering a cost-effective alternative with minimal modifications required for integration into existing SL-based models.

This scheme had not been considered in the original CATRINE proposal, however, the description and the results presented in Mortezaazadeh et al., 2024, and its suitability for the IFS, were motivating factors for a further deeper investigation. Use of a quadratic rather than cubic interpolation polynomial suggest another less obvious advantage: it is expected that the quadratic polynomial will produce smaller than the cubic polynomial overshoots/undershoots and therefore, as we saw in D1.1, the limiter will have smaller impact resulting in an overall more conserving scheme.

For this reason, we have implemented a test version of this scheme for the IFS model which has been tested in our plume case studies yielding very promising results as it can be seen in figure 8. Indeed, sweep interpolation is the most conserving, it reduces the conservation on average by 50% compared with cubic interpolation for LQM and LQM3D limiters. Combining it with LQM3DCONS produces the most conserving scheme which is as good or slightly better than a combination of COMAD with LQM3DCONS limiter. We note that at this stage the combination of COMAD with the sweep algorithm is not possible. However, this is an effort worth pursuing in WP2 of CATRINE as our current plots indicate that a combined sweep-COMAD scheme with LQM3DCONS limiter could reduce the conservation error even further. This is an attractive option not only for greenhouse tracers but for all the multi-tracer atmospheric composition forecasts run by CAMS as well, as it will improve both conservation properties and their computational efficiency.

It is also worth mentioning that so far, sweep has also been tested in the advection of water tracers in NWP forecasts yielding neutral results in terms of forecast accuracy confirming our expectations. We have also repeated the idealised tracer-correlation tests described in sec 5.1 confirming that it is equally accurate and preserves the same properties of the cubic interpolation when it is combined with any of the available limiters in IFS.



CATRINE

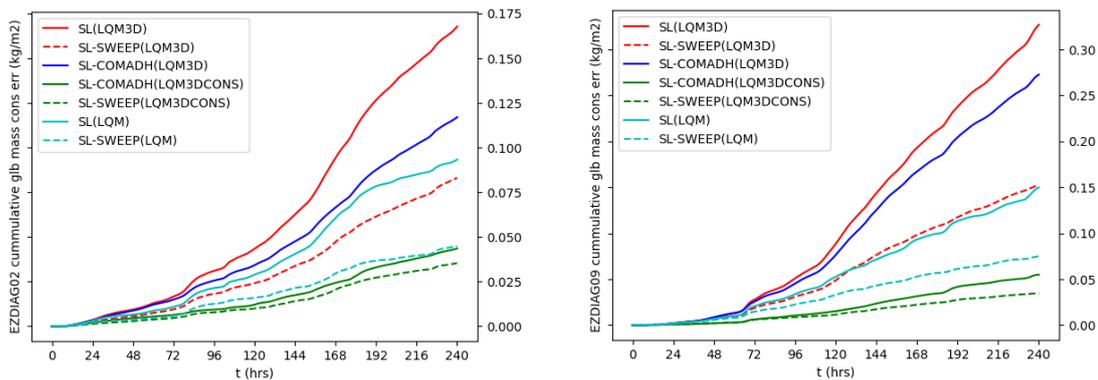


Figure 8. Accumulation of mass conservation error in kg/m² for same tracers (emitted from the near surface level) as in figure 5d. Here we compare sweep interpolation with the standard cubic interpolation scheme used in the IFS SL advection scheme (labelled SL) and with COMAD interpolation (labelled SL-COMADH) using LQM, LQM3D, LQM3DCONS limiters.

9 Concluding remarks and future work

In this deliverable we provide a detailed account of how the various items of Task 1.2 under WP1 of the CATRINE project have been addressed. These items are also summarised in section 2.2.2 of this report and focus on research aimed at developing an improved advection scheme for the global CO2MVS model IFS, specifically a scheme which has a larger degree of inherent conservation but maintains its good accuracy. We have conducted research in fundamental numerical aspects of transport schemes and, as expected by the nature of such work, some of our proposed ideas have yielded promising results while some others have proven to be less beneficial. Additionally, further work outlined in section 8 has been undertaken, based on very recent promising research published in the literature. The most successful investigations in this report, combined with findings from D1.1, have led us to the following conclusions:

1. Overall, the studies conducted in WP1 demonstrate that the mass fixer is performing very well, and that it eliminates the large sensitivity that has been found with respect to the application of different limiters. This has been confirmed by comparing different configurations of the advection scheme against observations. However, a greater degree of inherent local conservation is desirable for better consistency with governing equations, better preservation of background values and improved inversions. It is important that such improvements are not achieved at the expense of accuracy, as some options that enhance conservation can also be over-diffusive as we demonstrated in D1.1. This has been our main design objective for the improved scheme that we would like to deliver in WP2.
2. The above objectives for an improved scheme are satisfied largely by activating COMAD along the zonal and meridional directions for both convergent and divergent flows and combining it with the new LQM3DCONS monotone limiter and with the IFS mass fixer (Diamantakis and Agusti-Panareda, 2017) with a value of $\beta = 3$ to preserve better the background values. The mixed results with COMAD along the vertical show that this option is not yet ready for operational use. Work in this topic will conclude in T2.1 of WP2.
3. TL/AD code will be fully developed for the above scheme - this is already in progress and should be available during the initial stages of WP2.
4. The sweep interpolation combined with LQM3DCONS monotone limiter, and the mass fixer provides a possible alternative setup for CO2MVS. This will be further developed to include TL/AD. For optimal use in CO2MVS we aim to develop a combined scheme with COMAD which will bring greater reduction of mass conservation errors and limit the erroneous accumulation of mass in the boundary because of the SL formulation and its interaction with the surface boundary. This is an extra piece of development which has not been included in our deliverable and may be challenging to complete in WP2. We will try to advance it as much as possible and consolidate the work, perhaps post CATRINE if we are unable to complete it.
5. The alternative boundary condition described in sec 4.2 could provide a “fix” for the mass accumulation problem which can be observed near the surface under special meteorological conditions (winds converging toward a grid-point with calm conditions) enhancing COMAD and sweep interpolation schemes. Hence, more testing is needed to judge its suitability for our final scheme. There are other alternatives that can be pursued instead which will involve extension of the 3-point COMAD scheme described in section 3.4.1 and modifications already explored in section 4.3. Our final decision if any of these ideas could be used in practice and which one is the most appropriate for the IFS will be taken during the initial stages of WP2.

10 References

- R. Bermejo and J. Conde. A conservative Quasi-Monotone Semi- Lagrangian Scheme. Mon. Weather Rev., 130:423–430, 2002.
- R. Bermejo and A. Staniforth. The conversion of semi-lagrangian advection schemes to quasi-monotone schemes. Monthly Weather Review, 120(11):2622 – 2632, 1992.
- F. Chevallier, A. Agusti-Panareda, M. Krol, W. Peters, and S. Versick. D7.1 design of protocol for preliminary global model intercomparisons. CATRINE deliverable, 2024.
- M. Diamantakis, S. Malardel, L. Cantarello, C. Laurent, G. Tumolo, S. Versick and A. Agusti-Panareda. D1.1 Evaluation and analysis of IFS tracer advection scheme accuracy and conservation properties and its ability to estimate emissions, CATRINE deliverable, 2025.
- M. Diamantakis and A. Agusti-Panareda. A positive definite tracer mass fixer for high resolution weather and atmospheric composition forecasts. Technical Report 819, ECMWF, 2017.
- M. Diamantakis and J. Flemming. Global mass fixer algorithms for conservative tracer transport in the ECMWF model. Geosci. Model Dev., 7:965–979, 2014.
- M. Diamantakis and F. Vana. A fast converging and concise algorithm for computing the departure points in semi-Lagrangian weather and climate models. Q.J.R. Meteorol. Soc., 148(743):670–684, 2022.
- S. D. Eastham and D. J. Jacob. Limits on the ability of global eulerian models to resolve intercontinental transport of chemical plumes. Atmospheric Chemistry and Physics, 17(4):2543–2553, 2017. doi: 10.5194/acp-17-2543-2017. URL <https://acp.copernicus.org/articles/17/2543/2017/>.
- ECMWF. *IFS Documentation CY49R1 - Part III: Dynamics and Numerical Procedures*. <https://www.ecmwf.int/en/elibrary/81624-ifs-documentation-cy49r1-part-iii-dynamics-and-numerical-procedures>, 2024.
- M. Hortal. The development and testing of a new two-time-level semi-Lagrangian scheme (SETTLS) in the ECMWF forecast model. Q.J.R. Meteorol. Soc., 128:1671–1687, 2002.
- J. Kent, P.A. Ullrich and C. Jablonowski. Dynamical core model intercomparison project: Tracer transport test cases. Q.J.R. Meteorol. Soc., 140: 1279-1293, 2014.
- S. Malardel and D. Ricard. An alternative cell-averaged departure point reconstruction for pointwise semi-lagrangian transport schemes. Q.J.R. Meteorol. Soc., 2015.
- S. Malardel, N. Wedi, W. Deconinck, M. Diamantakis, C. Kuehnlein, G. Mozdzyński, M. Hamrud, and P. Smolarkiewicz. A new grid for the IFS. ECMWF Newsletter, No. 146-Winter 2015-16:30–36, 2016.
- Mortezazadeh, M., Cossette, J.-F., Dastoor, A., de Grandpré, J., Ivanova, I., and Qaddouri, A.: Sweep interpolation: a cost-effective semi-Lagrangian scheme in the Global Environmental Multiscale model, Geosci. Model Dev., 17, 335–346, <https://doi.org/10.5194/gmd-17-335-2024>, 2024.

CATRINE

F. Rabier, H. Jarvinen, E. Klinker, J.F. Mahfouf, and A. Simmons. The ECMWF operational implementation of four-dimensional variational assimilation. part i: experimental results with simplified physics. 126:1143–1170, 2000.

H. Ritchie, C. Temperton, A. Simmons, M. Hortal, T. Davies, D. Dent, and M. Hamrud. Implementation of the semi-Lagrangian method in a high-resolution version of the ECMWF forecast model. *Mon. Weather Rev.* 121, 489–514, 1995.

B. Sørensen, E. Kaas, and U. S. Korsholm. A mass-conserving and multi-tracer efficient transport scheme in the online integrated Enviro-HIRLAM model, 2013.

P.H. Lauritzen and J. Thuburn. Evaluating advection/transport schemes using interrelated tracers, scatter plots and numerical mixing diagnostics. *Q.J.R. Meteorol. Soc.*, 138: 906-918, 2012.

Willson, J. L., Reed, K. A., Jablonowski, C., Kent, J., Lauritzen, P. H., Nair, R., Taylor, M. A., Ullrich, P. A., Zarzycki, C. M., Hall, D. M., Dazlich, D., Heikes, R., Konor, C., Randall, D., Dubos, T., Meurdesoif, Y., Chen, X., Harris, L., Kühnlein, C., Lee, V., Qaddouri, A., Girard, C., Giorgetta, M., Reinert, D., Miura, H., Ohno, T., and Yoshida, R.: DCMIP2016: the tropical cyclone test case, *Geosci. Model Dev.*, 17, 2493–2507, <https://doi.org/10.5194/gmd-17-2493-2024>, 2024.

Document History

Version	Author(s)	Date	Changes
0.1 (Initial document created)	Sylvie Malardel, Michail Diamantakis	June 2025	Initial version
1.0	Sylvie Malardel, Michail Diamantakis	June 2025	Changes after internal reviews

Internal Review History

Internal Reviewers	Date	Comments
Anna Agusti-Panareda (ECMWF) and Stefan Versick (KIT)	June 2025	Internal review