Extending Legacy Climate Models by Adaptive Mesh Refinement for Single Component Tracer Transport

Yumeng Chen[1,2], Konrad Simon[1,2], Jörn Behrens[1,2]

[1]Department of Mathematics, Universität Hamburg, Hamburg, Germany

[2]Center for Earth System Research and Sustainability (CEN), Universität Hamburg, Grindelberg 5, 20144, Hamburg, Germany

Corresponding author: Yumeng Chen, yumeng.chen@uni-hamburg.de

Abstract

Integrating Adaptive Mesh Refinement (AMR) into climate models is problematic partly because several components have difficulty in accommodating adaptive grids. However, on coarse resolutions, errors from each component of climate models contribute to the overall errors of the model output. Using AMR in single components should reduce the overall model error. On the other hand, we can use AMR in existing climate models with significantly reduced development time compared to designing a new model equipped with AMR.

We integrate AMR into the tracer transport module of the atmospheric model ECHAM6 and test our implementation in several idealized scenarios on spherical geometries as well as in a realistic application scenario (dust transport). In order to achieve this goal, we modify the Flux-Form Semi-Lagrangian (FFSL) transport scheme in ECHAM6 such that we can use it on adaptive meshes while retaining all important properties such as mass conservation of the original FFSL implementation. Our proposed AMR scheme is dimensionally split and ensures that high-resolution information is always propagated on (locally) highly resolved meshes. We also introduce a data structure that can accommodate an adaptive Gaussian grid.

We demonstrate that our AMR scheme improves both accuracy and efficiency compared to the original FFSL scheme. More importantly, our approach improves the representation of transport processes in ECHAM6 for coarse resolution simulations. Hence, the results of this paper suggest that we can overcome the overhead of developing a fully adaptive earth system model by integrating AMR into single components while leaving data structures of the dynamical core untouched. This enables researches to retain well-tested and complex legacy code of existing models while still improving the model's accuracy.

Plain Language Summary

Mesh adaptivity is a valuable tool in many branches of computational sciences and can help to reduce the overall model error by only refining meshes in specific areas when actually necessary. Here we suggest a way to integrate mesh adaptivity into an existing earth system model, ECHAM6, without having to redesign the implementation from scratch. This is advantageous since, first, many effects can not be fully represented in long time simulations with standard meshes which do not permit a high resolution due to computational constraints and, secondly, since designing a fully adaptive model from scratch would be costly and time consuming.

Prototypically we show how to integrate adaptive meshes in a tracer transport module, i.e., a module in the earth system model that computes the evolution of certain substances (tracers) such as CO_2 or dust. We show that while the additional computational effort is manageable the error can be reduced compared to a low resolution standard model. Computational examples are presented for idealized test cases for which exact solutions are known and, prototypically, also for the evolution of Sahara dust as a real world scenario.

1 Introduction

The climate system is inherently multi-scale. In climate models, various processes are under-resolved because the resolution cannot represent details of these processes. One of the most straightforward approaches to better accuracy is increasing spatial resolution. However, high-resolution climate simulations are still computationally expensive, especially for long-term climate simulations like paleoclimate simulation. Adaptive Mesh Refinement (AMR) is an attractive alternative for global high-resolution climate models. The AMR technique refines and coarsens local meshes during run-time based on designated refinement criteria.

There is active research on AMR applications in the climate community dating back to the 1980s. For example, Skamarock and Klemp (1993) proposed an early non-hydrostatic model using AMR. More recently Jablonowski et al. (2009) constructed a finite volume general circulation model on a reduced lat-lon grid. Kopera and Giraldo (2015) constructed an atmospheric model using a Galerkin method on a cubed-sphere. These efforts focus on the dynamical cores of atmospheric models. Utilizing these methods for realistic climate simulations needs further research and development.

We propose an alternative pathway towards adaptivity in climate models to tackle concerns with AMR in operational climate models ranging from properties of numerical schemes to the coupling between dynamical core and physics packages (Weller et al., 2010).

Constructing a complete model from scratch usually takes decades of research. Instead, we propose to integrate AMR into single components of existing models, here ECHAM6, which could bring about immediate benefits. It is not uncommon to apply different resolutions for different components of a numerical model. For example, Berthet et al. (2019) showed that a high-resolution dynamical core using low-resolution parameterizations generates satisfactory results.

Enabling AMR in the passive tracer transport module of a climate model can improve the representation of the tracer transport process and it can potentially improve the general quality of climate simulations. The tracer transport module controls advective passive tracer transport processes in climate models. These tracers interact with wind in many other processes in the climate system and have feedback on the radiative balance or cloud formations. Consequently, these tracers affect the state of the climate system significantly.

Despite these benefits of integrating AMR into the tracer transport module of an existing model, there are still difficulties in achieving this goal:

- How does the tracer transport scheme handle hanging nodes on non-conforming adaptive meshes?
- How many improvements can we gain from integrating adaptive tracer transport schemes without refining other components?

We introduce AMR into the tracer transport module of ECHAM6. ECHAM6 uses the Flux-Form Semi-Lagrangian (FFSL) scheme (Lin & Rood, 1996). The scheme has two essential properties in climate models: mass conservation and semi-Lagrangian time stepping. Semi-Lagrangian schemes are particularly useful for the Gaussian grid in ECHAM6. The Gaussian grid is a variation of the lat-lon grid, where the longitude is equally spaced in the longitudinal dimension, and the latitude grid corresponds to Gaussian quadrature points for numerical integration. The Gaussian grid leads to smaller grid intervals around poles, which poses a CFL-limit on the time step size. Semi-Lagrangian time stepping ensures stable integration for large time steps.

However, on the adaptive mesh, the existing transport scheme in ECHAM6 cannot retain all its properties when hanging nodes are present. Hanging nodes lie at the interface between high-resolution and low-resolution areas. Ghost cells are a common treatment of hanging nodes. The scheme creates high-resolution ghost cells at low-resolution areas along the interface such that the stencil of the numerical scheme always lies at uniform resolutions. For example, Jablonowski et al. (2009) use ghost cells for the FFSL scheme but their implementation does not maintain the semi-Lagrangian time-stepping.

Another plausible approach is to substitute the existing transport scheme by a mass conservative semi-Lagrangian scheme, which can handle irregular meshes. For example, Nair and Machenhauer (2002) proposed a cell-integrated semi-Lagrangian scheme; Lauritzen et al. (2010) proposed a more efficient mass conservative semi-Lagrangian scheme using Stokes theorem. However, the comparison between the original climate model and the climate model with adaptive tracer transport would be difficult if we use two different transport schemes.

We propose a modified version of the existing tracer transport scheme which retains essential properties of the original scheme. Our modified tracer transport scheme allows us to reuse the code for vertical tracer transport and a class of limiters in the existing model without further investigation.

Utilizing idealized test cases, we quantitatively investigate the unique properties of our modified scheme on adaptive meshes and non-adaptive meshes even though many other tracer transport schemes using AMR are well studied (Behrens, 1996; Kessler, 1999; Iske & Käser, 2004; Jablonowski et al., 2006). In particular, we examine the effect of using coarse initial condition and wind field using idealized test cases as we only integrate AMR into a single component of the climate model.

We further validate our proposed AMR approach simulating the prototypical but realistic example of dust transport in ECHAM6. Dust is particularly suitable to demonstrate the effect of AMR since it has local sources and is transported around the entire globe. The global distribution of dust shows strong local features whose representation can get improved by local refinements.

The paper is organized as follows. We discuss our adaptive tracer transport scheme in Section 2. In order to quantitatively demonstrate the property of our modified scheme and features of AMR, we show our results in idealized tests in Section 3. We further demonstrate our idea of integrating AMR into tracer transport component of the existing model in Section 4 and conclude with a discussion of our results and future work in Section 5.

2 The Adaptive Transport Scheme

In order to ensure a fair examination of the partial introduction of AMR into the existing model ECHAM6, we use the original FFSL scheme in ECHAM6. The FFSL scheme is particularly suitable for climate models because it is accurate, efficient, mass conservative and semi-Lagrangian. The FFSL scheme is a combination of dimensionally split technique, 1-D finite volume transport scheme and Semi-Lagrangian extension for finite volume schemes.

Our aim is to use the FFSL scheme on adaptive meshes. However, we cannot extend the FFSL scheme to adaptive meshes while retaining all its properties without any modifications. We will explain the FFSL scheme, its problem on adaptive meshes and our modification in detail in this section.

2.1 The Flux-Form Semi-Lagrangian Scheme

We present the Flux-Form Semi-Lagrangian (FFSL) transport scheme proposed by Lin and Rood (1996). The FFSL scheme solves the 2-D transport equation. Climate models often rely on the transport equation in spherical coordinates:

$$\frac{\partial \rho c}{\partial t} + \frac{1}{a\cos\theta}\left(\frac{\partial \rho c u}{\partial \lambda} + \frac{\partial \rho c v \cos\theta}{\partial \theta}\right) = 0 \tag{1}$$

where a is the radius of the sphere, (λ, θ) is the longitude and latitude on the sphere, (u, v) is the horizontal velocity, ρ is the air density, c is the tracer concentration. For convenience of introducing the scheme, we set $c \equiv 1$.

The dimensionally split technique of the FFSL scheme is second order accurate in time. The method splits the 2-D transport equation in (1) into two 1-D transport equations:

$$\frac{\partial \rho}{\partial t} + \frac{\partial \rho u}{a\cos\theta\partial\lambda} = 0 \tag{2}$$

$$\frac{\partial \rho}{\partial t} + \frac{\partial \rho v\cos\theta}{a\cos\theta\partial\theta} = 0 \tag{3}$$

The dimensionally split technique eases the difficulty in extending 1-D methods into higher dimensions and enables the application of various 1-D limiters to 2-D problems.

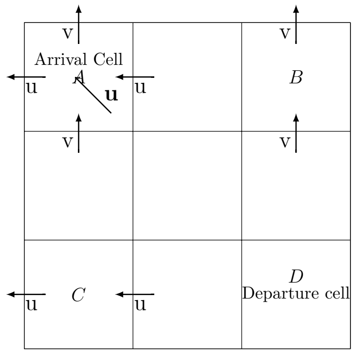

Figure 1. Schematic illustration of the dimensionally split scheme. If the arrival cell is cell A and departure cell is cell D, the dimensionally split scheme transports information from the cell D to cells B and C respectively. Then, the scheme updates the value in cell A from cells B and C.

This method is equivalent to the COSMIC splitting proposed in Leonard et al. (1996). The advantage of the FFSL scheme is that the scheme leads to a mass conservative and consistent dimensionally split technique since the Strang splitting cannot preserve both mass conservation and consistency condition for tracer transport problems.

The FFSL scheme defines a 1-D conservative operator for the flux difference of two cell edges $F_C(\rho)$:

$$F_C^{\lambda}(\rho) = \frac{\partial \rho u}{a \cos\theta \partial \lambda} \qquad\qquad F_C^{\theta}(\rho) = \frac{\partial \rho v \cos\theta}{a \cos\theta \partial \theta} \tag{4}$$

Here, the subscript C means the operator is conservative and the superscript represents the dimension of the 1-D operator. The dimensionally split technique allows any 1-D finite volume transport scheme to solve the 1-D operator $F_C(\rho)$. The finite volume scheme ensures the mass conservation of the FFSL scheme.

In order to achieve the consistency condition of the FFSL scheme, the scheme also uses an advective operator, which is a variation of the $F_C(\rho)$:

$$F_A^{\lambda}(\rho) = F_C^{\lambda}(\rho) - \Delta t \rho(\nabla \cdot u) \qquad\qquad F_A^{\theta}(\rho) = F_C^{\theta}(\rho) - \Delta t \rho(\nabla \cdot u) \tag{5}$$

where A means the operator only solves the advective part of the transport equation, Δt is the time interval and $\nabla \cdot u$ is the divergence. The second term of equation (5) is computed by a 2^{nd} order finite difference scheme.

Similar to the Strang splitting, the FFSL scheme alternates the direction sequentially. The dimensionally split scheme first solves the 1-D equation in λ or θ dimension.

$$\rho_A(\lambda) = \rho^n + F_A^\lambda(\rho^n) \qquad\qquad \rho_A(\theta) = \rho^n + F_A^\theta(\rho^n) \tag{6}$$

the superscript n is the current time step. The scheme uses the advective operator $F_C(\rho)$ as the inner operator, which guarantees the consistency condition.

Using ρ_A as the initial condition, the scheme subsequently solves the 1-D equation in the other direction.

$$\begin{aligned}
\rho(\rho_A(\lambda), \rho^n) &= \rho^n + F_C^\lambda(\rho^n) + F_C^\theta(\rho_A(\lambda)) \\
\rho(\rho_A(\theta), \rho^n) &= \rho^n + F_C^\theta(\rho^n) + F_C^\lambda(\rho_A(\theta))
\end{aligned} \tag{7}$$

The mass conservation is guaranteed by the conservative outer operator. Results of $\rho(\rho_A(\lambda), \rho^n)$ and $\rho(\rho_A(\theta), \rho^n)$ tilt to different directions. Hence, the final solution for the next time step, $n+1$ is the average of the outer operator in each direction:

$$\rho^{n+1} = \frac{1}{2}(\rho(\rho_A(\lambda), \rho^n) + \rho(\rho_A(\theta), \rho^n)) \tag{8}$$

We illustrate the scheme in figure 1. If the cell D is the departure cell corresponding to cell A, the scheme transports information dimensionally from cell D to cells B and C, which in turn are the departure cells of cell A in each dimension. Therefore, the value of the arrival cell A is calculated based on cells B and C.

2.2 Semi-Lagrangian Extension on Adaptive Meshes

The FFSL scheme attains long time steps by a semi-Lagrangian extension from 1-D finite volume schemes (Leonard et al., 1995). Similar to traditional semi-Lagrangian schemes, the extension requires computation of trajectories described by the flow field. However, by construction, the extension also requires the mass flux of each cell edge during one time step, which is a sweep of mass along trajectories. This semi-Lagrangian computation takes account for the exact integration of mass flux across an edge, similar to a finite volume scheme, and thus yields mass conservation. In order to improve the efficiency of the implementation, the FFSL scheme employs the widely used idea of cumulative mass first described in Colella and Woodward (1984). The cumulative mass of a cell is the mass from the beginning of the domain to the cell. Thus, the mass along the trajectory is the difference between the arrival cell and the departure cell, and the finite

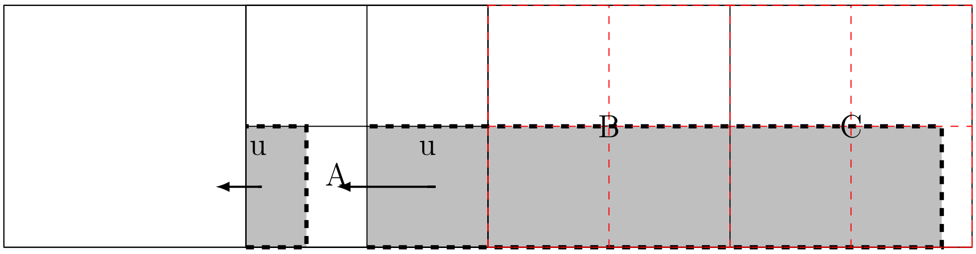

Figure 2. Illustration of the semi-Lagrangian extension for finite volume schemes on adaptive meshes. cell A is the arrival cells. The dashed read cells are ghost cells. The shaded areas represent departure trajectories, which is the mass flux at the edge of the arrival cell.

volume flux at the departure cell. The cumulative mass significantly reduces the computational cost.

However, when using the semi-Lagrangian extension on adaptive meshes, problems arise. The FFSL scheme assumes a structured rectangular grid, where the cell centers align with each other in each dimension such that the dimensionally split scheme can use 1-D solvers for each dimension. For example, the cell center always lies at the same latitude when the scheme computes for longitudinal direction. However, hanging nodes on adaptive meshes cannot guarantee an alignment as shown in Figure 2. Breaking the alignment assumption leads to inconsistency and violates mass conservation. For example, if a 1-D finite volume scheme computes the value of the next time step at the arrival cell A in Figure 2, the 1-D scheme could include the mass at the entire cell B while a consistent treatment needs only the mass at the lower shaded area of cell B.

In order to satisfy the alignment assumption, we can use ghost cells, which are the red cells in Figure 2. However, using ghost cells for large Courant numbers prevents the scheme from using cumulative mass since it is difficult to define the cumulative mass for high-resolution cells. Without cumulative mass, the semi-Lagrangian extension may lead to multiple computations of the mass because the departure trajectory of different edges may overlap, leading to an inefficient scheme.

2.3 Modified Flux-Form Semi-Lagrangian Scheme

As described in Section 2.2, the original FFSL scheme cannot handle hanging nodes efficiently because it uses a finite volume scheme with a semi-Lagrangian extension to

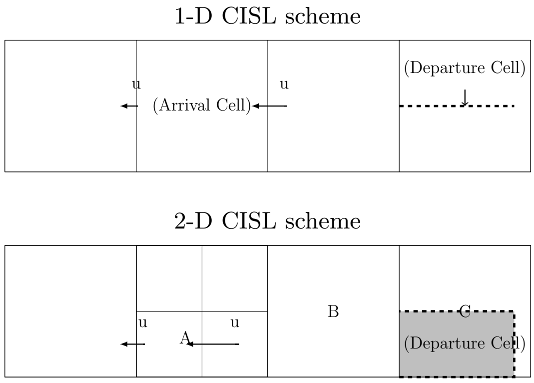

Figure 3. Illustration of the CISL scheme in 1-D and 2-D settings, where the 1-D setting does not account for the variation in two dimensions and the 2-D setting accounts for the variation in two dimensions.

solve 1-D problems, where it is computationally expensive to obtain the mass along the trajectory. We expect that a mass conservative semi-Lagrangian scheme without the sweep along trajectories can solve the problem arising with hanging nodes. The cell-integrated semi-Lagrangian (CISL) scheme (Nair & Machenhauer, 2002) is a good candidate. Instead of adding up the mass along the whole trajectory of cell edges, the CISL scheme update values from the mass at departure cells. In particular, Lauritzen (2007) shows that the CISL scheme is an alternative point of view of Godunov-type finite volume schemes with a semi-Lagrangian extension. Hence, we can safely substitute the finite volume scheme to the CISL scheme and expect a fair comparison of numerical results on adaptive and non-adaptive meshes.

Similar to finite volume schemes, in a 1-D setting, the CISL scheme assumes the cell center value as the cell average:

$$\rho_i^c = \frac{1}{\Delta x_i} \iint_{\Delta x_i} \rho dx \tag{9}$$

where x is either λ or $\mu = \sin\theta$ and Δx_i is the interval of a cell i. The integrand is a sub-cell reconstruction function based on the cell center value. For example, the Godunov scheme assumes the sub-cell reconstruction function as constant.

In the CISL scheme, the departure cell is formed by the departure position of the cell edges of the arrival cell and the 1-D scheme updates values from the departure cell:

$$\rho_i^{n+1}(\lambda) = \frac{1}{\Delta\lambda_i} \int_{\Delta\lambda_d} \rho^n d\lambda \qquad\qquad \rho_i^{n+1}(\theta) = \frac{1}{\Delta\mu_i} \int_{\Delta\mu_d} \rho^n d\mu \tag{10}$$

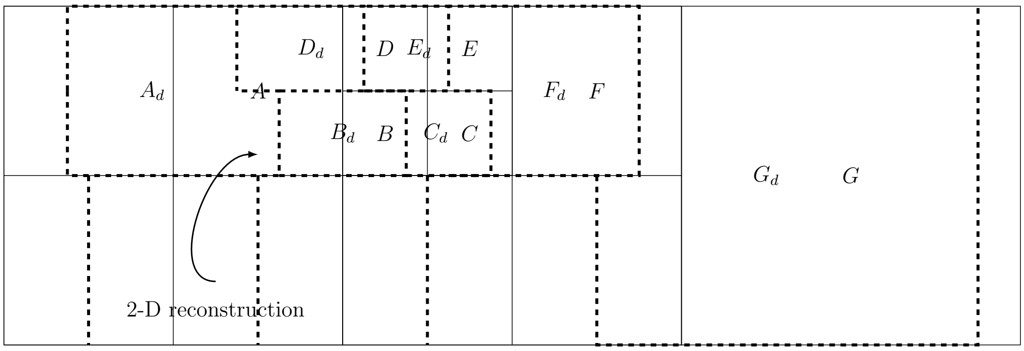

Figure 4. Illustration of the stability issue of hanging nodes. The solid mesh is the underlying Eulerian grid while the dashed mesh is the Lagrangian mesh at previous time step. The dashed mesh is marked by subscript d.

where $\Delta\lambda_i$ and $\Delta\mu_i$ represent the interval of arrival cells in each dimension, $\Delta\lambda_d$ and $\Delta\mu_d$ are interval of departure cells in each dimension. As shown in Figure 3, the dashed line is the departure cell in 1-D. The scheme gets new values from the mass at the departure cells, which is an integral of the sub-cell reconstruction function over the interval of departure cells. The departure position of cell edges in each dimension on the sphere is described by:

$$\frac{a\cos\theta d\lambda}{dt} = u \qquad \frac{ad\mu}{dt} = v\cos\theta \tag{11}$$

Here, we use a first-order Euler method to solve the ODE as done in ECHAM6. The CISL scheme avoids the computation of mass along the trajectory while keeping the advantage of long time steps on adaptive meshes.

On an adaptive mesh with hanging nodes, the 1-D integral in Equation (10) does not consider the mass variation in the other dimension, which breaks the 2-D mass conservation. Therefore, we must use a 2-D integral:

$$\rho^{n+1}(\lambda) = \frac{1}{\Delta A_i}\iint_{\Delta A_{\lambda,d}} \rho^n d\lambda d\mu \qquad \rho^{n+1}(\theta) = \frac{1}{\Delta A_i}\iint_{\Delta A_{\theta,d}} \rho^n d\lambda d\mu \tag{12}$$

where ΔA_i is the area of the arrival cell, $\Delta A_{\lambda,d}$ is the area of the departure cell in λ direction and $\Delta A_{\theta,d}$ is the area of the departure cell in θ direction. The cell interval in Equation (10) is different from the area in Equation (12) since the cell interval only considers a 1-D problem while the area considers the variation in 2-D as shown in Figure 3. By definition of the cell center value in Equation (9), Equation (12) can be reduced to

Equation (10) if departure cell and underlying Eulerian cell have the same refinement level.

The equivalence between Equations (10) and (12) allows us to use 1-D and 2-D reconstructions for different conditions. As shown in Figure 4, we apply a 2-D reconstruction function on adaptive meshes when a departure cell has a lower refinement level than the arrival cell. Otherwise, we apply a 1-D reconstruction function. For example, in Figure 4, cell D needs only a 1-D reconstruction function since it contains the departure cell at the same size as itself while cell A requires a 2-D reconstruction function since the departure cells located at cell A has different sizes.

In order to be consistent with the original implementation, we choose the same reconstruction function as the one used by the FFSL scheme in ECHAM6 such that we can make a fair comparison between the AMR scheme and the original scheme in the following sections and our idealized tests can provide insight for realistic simulations. The default option of the FFSL scheme in ECHAM6 uses the Piecewise Parabolic Method (PPM) as 1-D finite volume solver. The PPM is a finite volume Godunov-type method, which assumes a quadratic subcell distribution function. Interested readers can refer to Colella and Woodward (1984) for a detailed description of the PPM. We use a 1-D second order polynomial and a quasi-2D reconstruction as in Nair and Machenhauer (2002):

$$\rho(\lambda, \mu) = \begin{cases} \rho^c + \delta a^x x^2 + b^x(\frac{1}{12} - x^2) & l_d >= l \\ \rho^c + \delta a^\lambda \lambda^2 + b^\lambda(\frac{1}{12} - \lambda^2) + \delta a^\mu \mu^2 + b^\mu(\frac{1}{12} - \mu^2) & l_d < l \end{cases} \tag{13}$$

where $x \in (-\frac{1}{2}, \frac{1}{2})$ is either λ or μ in 1-D case, the condition l represents the refinement level of the Eulerian cell, l_d represents the refinement level of the departure cell, the coefficients a and b are computed following Colella and Woodward (1984). Because a and b are computed by 1-D interpolations, we remap the coarse cell values to refined cells by recursively using Equation (13) to form the interpolation stencil. The 2-D reconstruction function can also be used in the fully 2-D schemes as in the original work of Nair and Machenhauer (2002). The dimensionally split scheme benefits from the simplicity of the implementation in that the computation of the departure cell's position is still 1-D and the departure cell's shape is more regular than in a fully 2-D scheme.

Using our modified 1-D operator in the FFSL scheme, the original $F_C^d(\rho)$ in Section 2.1 becomes:

$$F_C^\lambda(\rho) = \rho^{n+1}(\lambda) - \rho^n \qquad\qquad F_C^\theta(\rho) = \rho^{n+1}(\theta) - \rho^n \tag{14}$$

where ρ is the updated value in Equation (10).

Our modified operator for the dimensionally split scheme retains the semi-Lagrangian time stepping. Moreover, the efficiency of the CISL scheme is similar to the original finite volume scheme with a semi-Lagrangian extension. Finally, the scheme is mass conserving as is the original scheme.

3 Idealized Tests

We implement the AMR scheme based on the data structure from Chen et al. (2018). There are a number of necessary considerations when using AMR, including errors arising from the AMR procedure or the choice of refinement criteria and their corresponding thresholds. Idealized tests can expose the accuracy and efficiency of the AMR scheme under various conditions. We can even design our experiments using idealized tests to mimic the behavior of our intended application since we plan to integrate the adaptive tracer transport scheme into an existing model while keeping other components unchanged.

We conduct idealized tests to demonstrate three essential aspects of our AMR scheme. Firstly, we show that our dimensionally split AMR scheme needs a special treatment as refinement strategy. Secondly, we examine various properties of our AMR scheme, including accuracy, efficiency and mass conservation. Thirdly, we explore the accuracy of the solution on adaptive meshes in situations where the AMR scheme interpolates low-resolution wind fields to high-resolution meshes.

We utilize three test cases: a solid body rotation test case (Williamson et al., 1992), a divergent test case (Nair & Lauritzen, 2010) and a moving vortices test case (Nair & Jablonowski, 2008). Each test case poses different challenges to our transport scheme. Hence, we can demonstrate that our AMR scheme possesses all numerical properties essential for our purpose.

The solid body rotation test case has a discretely divergence-free wind field and in the theoretical absence of diffusion the shape of the tracer distribution should not change during run-time. In the solid body rotation test case, the flow orientation can be controlled by the parameter α, where α is the angle between the flow orientation and the equator. This test case is challenging when the tracer moves around the poles due to the convergence of coordinate lines. It is a useful test case to explore accuracy and efficiency of our numerical scheme under idealized circumstances.

The divergent test case deforms the tracer distribution with a divergent wind field. Divergent wind is especially challenging for large time steps since the transport scheme needs to correctly move the tracer when the divergent wind leads to a high gradient in the tracer concentration.

Different from the solid body rotation test case and the divergent test case, the moving vortices test case distributes tracer over the entire globe. The moving vortices test case also severely deforms the tracer and the vortices form filaments in the tracer concentration. Strong deformation leads to discontinuities and, furthermore, it poses challenges for the AMR scheme because improper refinement criteria may result in refinement of the entire domain.

In these idealized tests, we measure the numerical results quantitatively in the ℓ_2 and ℓ_∞ error norms:

$$\ell_2 = \frac{\sqrt{\sum_i^{n_t^{cell}} (q_i - q_i^{exact})^2 dA_i}}{\sqrt{\sum_i^{n_t^{cell}} (q_i^{exact})^2 dA_i}} \tag{15}$$

$$\ell_\infty = \frac{\max |q_i - q_i^{exact}|}{\max |q_i^{exact}|} \tag{16}$$

where q_i is the tracer concentration in the ith cell, q_i^{exact} is the exact solution in the ith cell and dA_i is the cell area of the ith cell. In order to test the performance of our AMR scheme, we do not apply any limiters to the scheme in idealized tests.

In many tests, we need to investigate the number of cells in a simulation. The number of cells changes with time on adaptive meshes. In order to show the overall number of cells at each test, we average the number of cells over time:

$$\text{cell number} = \sum_t^{nt} \frac{n_t^{cell}}{nt} \tag{17}$$

where nt is the number of time steps, n_t^{cell} is the number of cells at time step t. The cell number can effectively and objectively reflect the efficiency of the AMR scheme regardless of the optimizations applied to the rest of the code.

3.1 Grid Refinement for Intermediate Steps

The dimensionally split scheme differs from genuinely multi-dimensional schemes creating a need for different refinement strategies.

Multi-dimensional schemes mimic the behavior of the multi-dimensional transport equation. These schemes get information at the new time step directly from the departure point along the trajectory. AMR schemes refine the departure areas and the arrival areas, and hence information always resides on the fine-resolution mesh.

The dimensionally split scheme also gets information from the departure point. However, as indicated in Figure 1, the scheme moves the information from the departure point to intermediate positions before moving the information to the arrival point. AMR schemes need to track this information and need to refine intermediate steps.

Using the solid body rotation test case as an example, we compare numerical errors between two refinement strategies. One strategy refines intermediate steps whereas the other does not refine intermediate steps. The flow transports the tracer around the globe with an angle of $\alpha = 0$ and $\alpha = \frac{3\pi}{20}$ with respect to the equator. These two settings lead to different maximum Courant numbers $\frac{\mathbf{u}}{\Delta x}\Delta t$, which shows the speed of information propagation in one time step. Here, $\mathbf{u}$ is the wind speed, Δx is the grid space, and Δt is the time step size.

In dimensionally split schemes, large Courant numbers can highlight the displacement between intermediate steps and final results because the information propagation is far away from the departure cell. When $\alpha = 0$, there is no divergence in each dimension in the wind field and the AMR scheme allows arbitrarily large Courant numbers. We use a Courant number of around 6 over the globe.

The dimensionally split scheme poses a limit to the Courant number as the dimensionally split scheme essentially performs 1-D semi-Lagrangian schemes. The divergence-free wind field in 2-D can be a result of the cancellation of 1-D divergence wind. When $\alpha = \frac{3\pi}{20}$, we use a maximum Courant number around 12 in the longitudinal direction, which is the largest Courant number without the crossing of trajectories in 1-D. We note that the transport of cosine bell is only affected by local small Courant numbers and extremely large Courant numbers can only occur around poles.

In order to expose the difference in these two refinement strategies, we use different spatial resolutions and keep the Courant number roughly fixed. Note that the Courant number is not exactly the same on different resolutions as the grid spacing changes with the latitude. The AMR scheme uses a gradient-based refinement criterion.

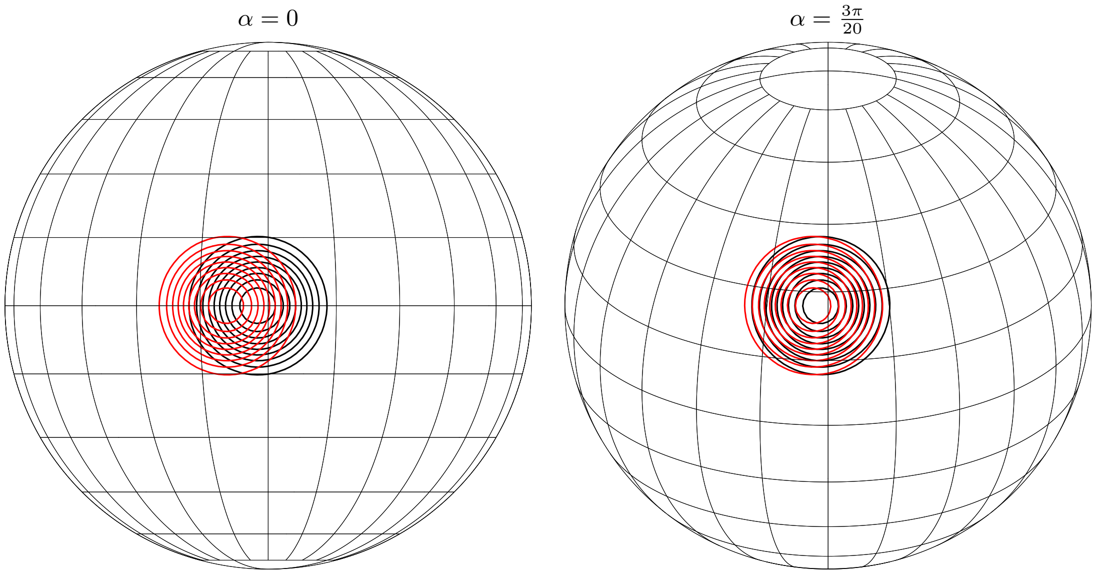

Figure 5. Illustration of the displacement of the numerical solution between the intermediate step after update in latitudinal direction and final results. The red distribution is the intermediate step and the black distribution is the final result. The initial condition is on the mesh for the previous time step. While the intermediate step and final results are on the mesh for the new time step. When $\alpha = 0$, the flow orientation is parallel to the equator and the Courant number is around 6. When $\alpha = \frac{3\pi}{20}$, the angle between the flow and the equator is $\frac{3\pi}{20}$ and the maximum Courant number of roughly 12 in the longitudinal direction occurs around poles while the Courant number around cosbells is much smaller than around poles.

When $\alpha = \frac{3\pi}{20}$, the threshold for mesh refinement is $\theta_r = 10^{-3}$ and the threshold for coarsening is $\theta_c = 5 \times 10^{-3}$. When $\alpha = 0$, $\theta_r = 5 \times 10^{-6}$ and $\theta_c = 5 \times 10^{-5}$.

In Figure 5, we illustrate how both flow orientations induce displacements between intermediate steps and final results under both flow orientations on a mesh with $1.25° \times 1.25°$ spatial resolution. The displacement is more visible when the tracer rotates along the equators due to different Courant numbers.

Figure 6 shows the numerical errors of these two refinement strategies. When $\alpha = \frac{3\pi}{20}$, numerical errors and the convergence rate of these two refinement strategies are comparable. Similar results arise from small displacements between intermediate steps and final results as shown in Figure 5. Our local high-resolution areas cover intermediate steps due to our sensitive refinement criterion.

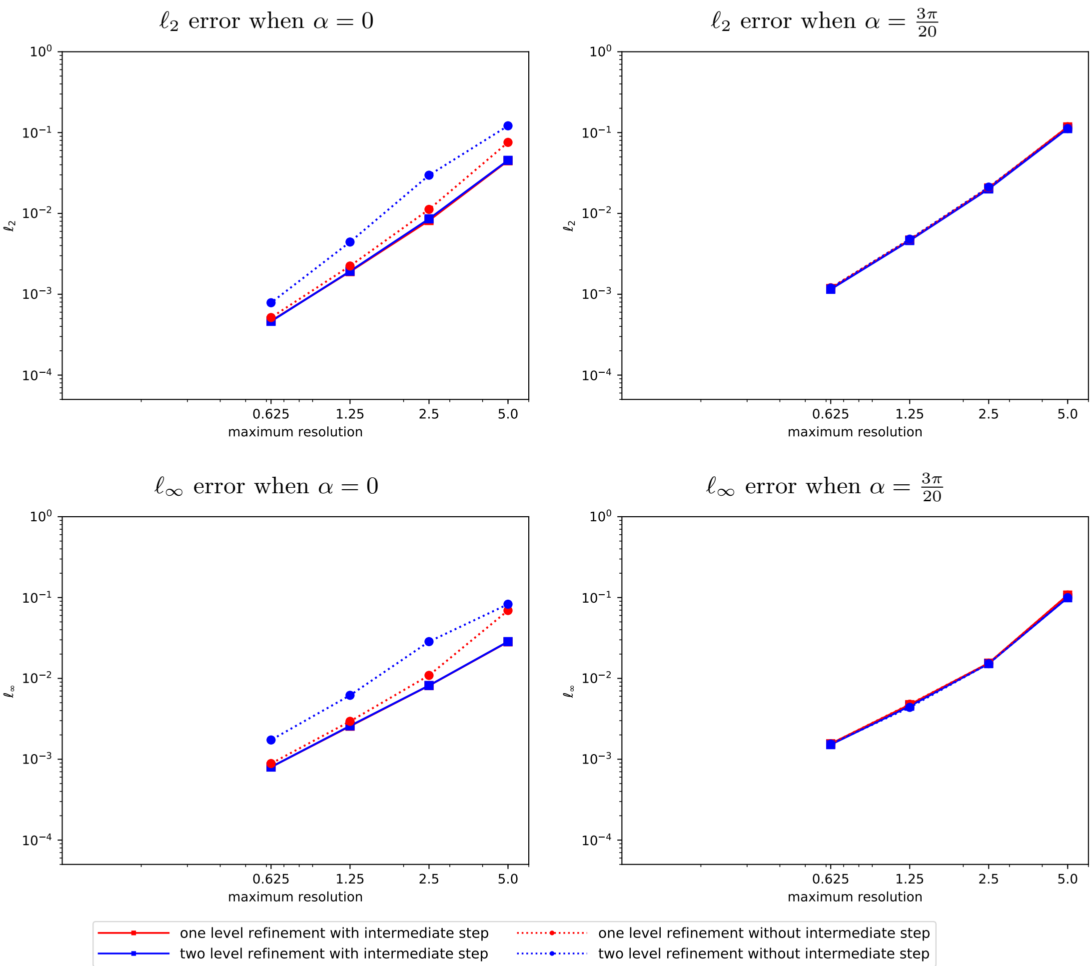

Figure 6. Comparison of the error of the solid body rotation test case after 12 days between refinement with intermediate step and refinement without intermediate step. Solid lines with rectangular markers show results with refinement at intermediate steps and dashed lines with circles show results without refinement at intermediate steps.

Numerical errors show a significant difference between these two refinement strategies when $\alpha = 0$. Without refining intermediate steps, the numerical error is higher on adaptive meshes than on non-adaptive meshes due to the coarse resolution. The AMR scheme leads to similar accuracy on adaptive meshes and non-adaptive meshes when the numerical scheme refines intermediate steps. Our implementation exposes the difference as the AMR scheme transports information from the mesh for the previous time step to the mesh for the new time step. Computations for both intermediate and final time step exist on the mesh for the new time step.

Our results demonstrate the schematic illustration of the dimensionally split scheme. The refinement of intermediate steps is essential for better accuracy when the Courant number is large. Although it is unlikely that the numerical model uses an extremely large Courant number away from the poles, we refine intermediate steps to ensure the accuracy.

3.2 Numerical Accuracy and Efficiency

The transport scheme behaves differently under different initial conditions and flow features. We examine the accuracy, efficiency and mass conservation of our AMR scheme using three different test cases.

3.2.1 Non-Divergent Flow with Local Tracer Distribution: The Solid Body Rotation Test Case

We examine our adaptive transport scheme in the solid body rotation test case. The solid body rotation test case has discretely non-divergent flow. The non-divergent flow also does not severely distort the tracer distribution and the gradient of the tracer does not change during the test. Hence, we can test the numerical properties in an ideal condition.

The test case uses a local tracer distribution with a radius of a third of the earth's radius. The test case allows us to initialize the tracer distribution on high-resolution adaptive meshes. The AMR scheme should result in very local high-resolution areas.

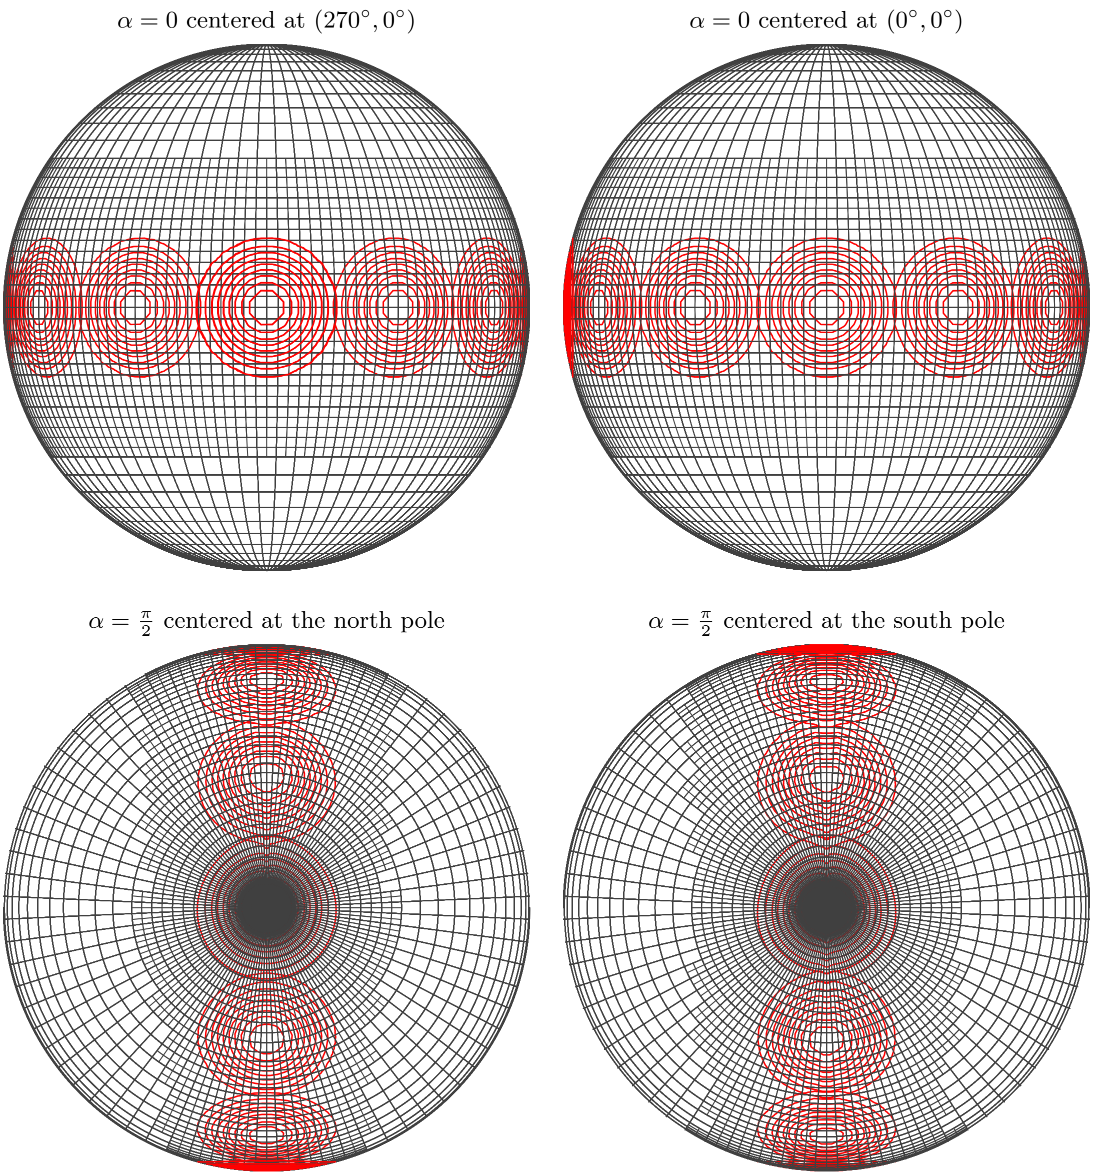

Figure 7. Snapshots of the solid body rotation test case when $\alpha = 0$ and $\alpha = 0.5\pi$ at each day with one level refinement. The coarse mesh has a resolution of $5° \times 5°$ and high resolution areas have a resolution of $2.5° \times 2.5°$.

We set the flow orientation $\alpha = 0$ and $\alpha = \frac{\pi}{2}$. When $\alpha = 0$, the tracer rotates around the globe parallel to the equator. When $\alpha = \frac{\pi}{2}$, the flow leads to cross-pole transport, which suffers from the geometrical problem of Gaussian grids at poles.

We test these two flow orientations with a maximum Courant number around 1 and 6 respectively. The AMR scheme utilizes a gradient-based criterion. Our threshold for cell refinement is $\theta_r = 0.02$ and the threshold for cell coarsening is $\theta_r = 0.015$ when $\alpha = \frac{\pi}{2}$ while the threshold for $\alpha = 0$ is the same as in Section 3.1.

As shown in Figure 7, the cosine bell is located in the high-resolution areas throughout the simulation, which exhibits that the refinement criterion detects regions where the gradient is present. The large high-resolution areas are also a result of the refinement strategy, where the intermediate steps are refined.

The distribution of the mesh explains the numerical accuracy of our transport scheme on adaptive meshes. The discrete representation of the tracer is similar on high-resolution areas of adaptive meshes and on the uniformly refined grid if both grids have the same maximum resolution. Figure 8 shows that the accuracy on adaptive meshes and non-adaptive meshes is similar when the scheme runs with the same maximum resolution on the mesh, which leads to similar convergence rates on these meshes.

Figure 8 also shows that the AMR scheme demands fewer cells than non-adaptive schemes to achieve similar accuracy. We also note that higher-order refinement does not necessarily result in fewer cells on the mesh.

The Gaussian grid contains a higher number of cells around the poles than elsewhere. Figure 9 shows that the AMR scheme captures the changing number of grid cells at different latitudes, especially the peak of cell number when tracer is located at poles.

We explore numerical accuracy, efficiency, and the convergence rate of the adaptive transport schemes in an ideal context, where we use a high-resolution initial condition and a non-divergent wind field. Our adaptive transport scheme can achieve similar accuracy to the scheme on non-adaptive meshes using reduced number of cells.

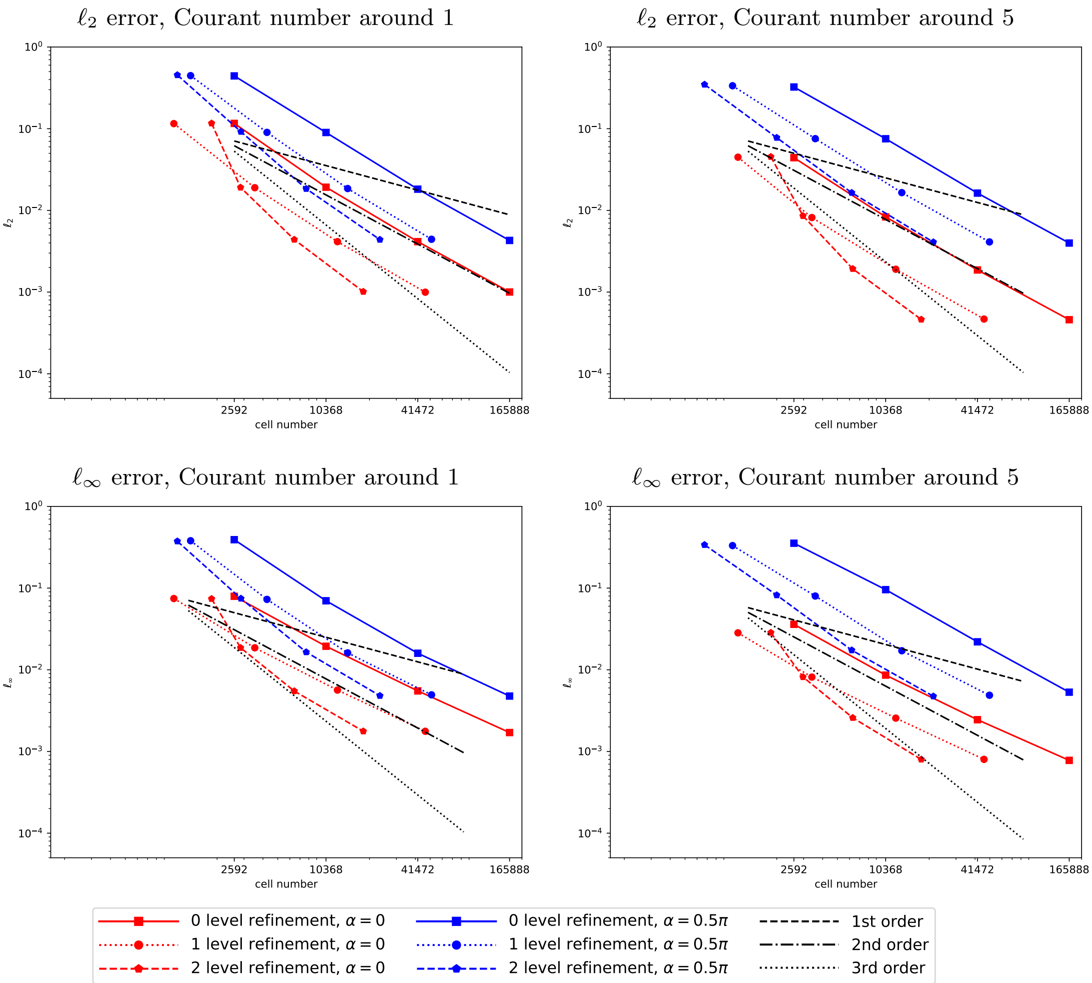

Figure 8. Convergence rate of the numerical results with respect to the number of cells in the solid body rotation test case. The red lines indicate the numerical results of tracer rotating along the equator while the blue lines are results of tracer rotating cross the poles.

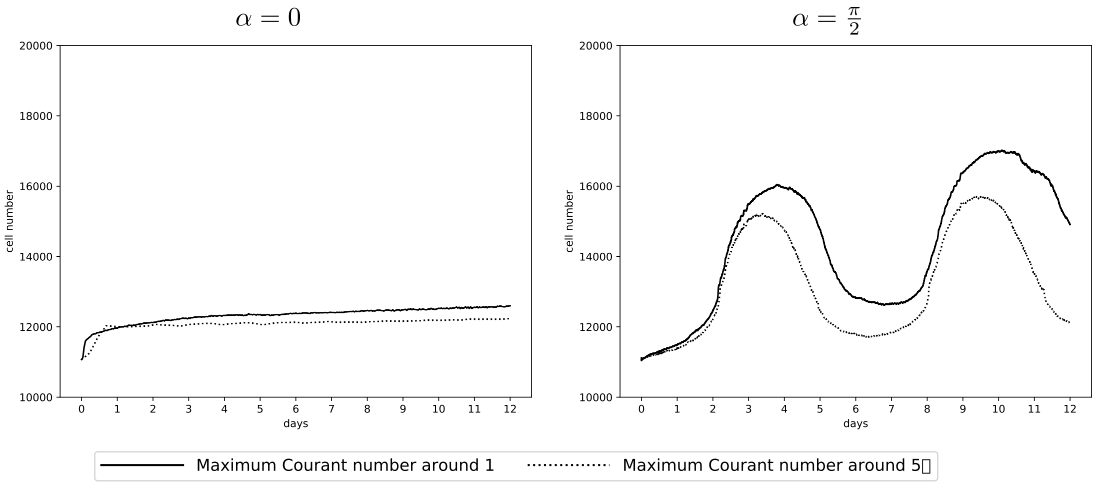

Figure 9. Evolution of cell number for rotation around the equator (in the left) and cross-pole transport (in the right) in the solid body rotation test case with a resolution of $2.5^\circ \times 2.5^\circ$. The solid line shows the cell number evolution with time when the Courant number is small and the dashed line show the cell number evolution with time when the Courant number is large.

3.2.2 Divergent Flow with Local Tracer Distribution: The Divergent Test Case

We test our AMR scheme in the divergent test case. The magnitude and the direction of the wind change swiftly in a divergent flow. The swift change of wind challenges the accuracy of our semi-Lagrangian scheme, which needs the correct departure position. Furthermore it may reveal inexact mass conservation, since concentration values will change to compensate for converging or diverging trajectories.

In the divergent test case, background flow transports two cosine bells along the equator while the divergent flow stretches them. From day 6 on, the test case reverses its flow and the tracer restores to its initial state. The final tracer distribution at day 12 is the same as the initial condition. There is no analytical solution for the test case but we can compare the final state with the initial condition to get the numerical error of the results.

Similar to the solid body rotation test case, the tracer distribution does not cover the entire domain but locates at limited areas. However, the size of the tracer is larger in the divergent test case than in the solid body rotation test case. The AMR scheme might need more grid cells to cover the whole tracer. We initialize the tracer distribution on the high-resolution areas and use a gradient-based refinement criterion. Our threshold for the refinement is $\theta_r = 0.2$ and the threshold for the coarsening is $\theta_c = 0.15$.

In the divergent test case, we take three steps to verify the performance of our AMR scheme. We first run the test case with and without one level refinement using a Courant number around 1 using a resolution of $5° \times 5°$ and demonstrate the representation of the tracer on high-resolution mesh.

As shown in Figure 10, the refinement criterion captures the distribution of the tracer completely. As the tracer gets stretched during the runtime, the high-resolution area leads to a better representation of filaments. Using high-resolution meshes significantly improves the representation of the tracer compared to using low-resolution meshes. The final tracer distribution is not completely the same as the initial condition as the divergent flow leads to a damping and distortion in numerical results.

Secondly, we use multiple levels of refinement to verify the sensitivity of the refinement level to the numerical accuracy and efficiency. The AMR scheme runs with an initial resolution of $20° \times 20°$. The refinement on adaptive meshes ranges from two level refinement up to 5 level refinement resulting in a resolution up to $0.625° \times 0.625°$ using a Courant number around 5.

As shown in Figure 11, we observe a similar convergence rate between uniformly refined meshes and locally refined meshes. Our results show that our AMR scheme and the non-AMR scheme generate numerical results with similar accuracy where the AMR scheme requires only a reduced number of cells in the divergent flow.

At last, we inspect another aspect of numerical accuracy: mass conservation. We show the evolution of relative mass in the divergent test case when the maximum resolution is $0.625° \times 0.625°$ with no adaptive refinement and one level refinement with a coarse resolution of $1.25° \times 1.25°$. We define the relative mass:

$$\text{relative mass} = \frac{\text{mass} - \text{mass}_{\text{mean}}}{\text{mass}_{\text{mean}}} \tag{18}$$

where mass is the mass at individual time step and $\text{mass}_{\text{mean}}$ is the temporal average of the mass in all time steps.

We observe that mass is conserved without AMR in Figure 12. However, mass declines with AMR experiments. After 960 time steps, the loss of relative mass is at an order of 10^{-12}. The loss of mass arises from the accumulation of rounding error of floating-

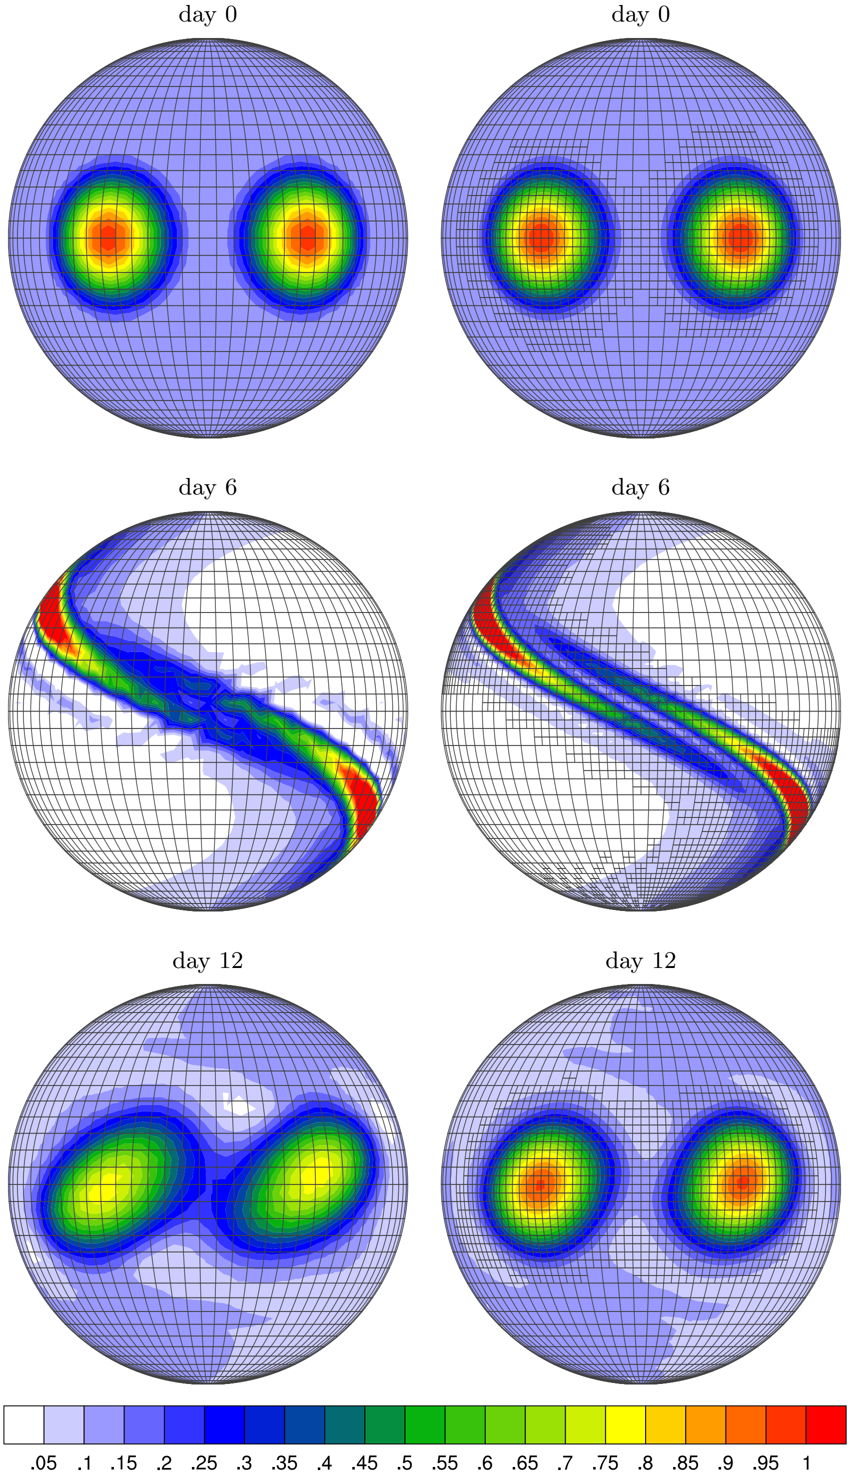

Figure 10. Numerical results of the divergence test case with a resolution of $5^\circ \times 5^\circ$ on the left panel and one level refinement on the right panel. The maximum resolution is $2.5^\circ \times 2.5^\circ$. The Courant number is around 1.

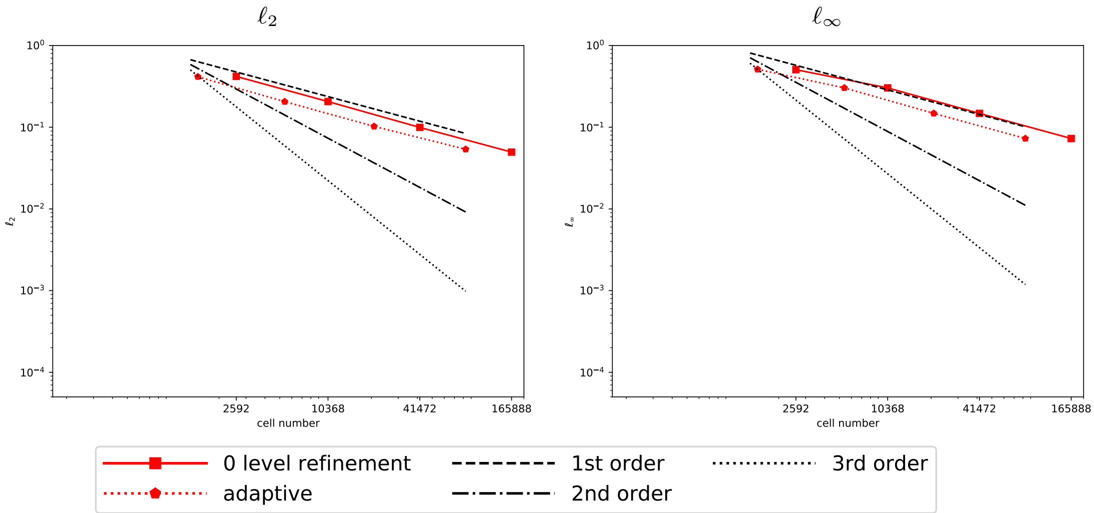

Figure 11. Convergence rate of the numerical results with respect to the number of cells in the divergent test case using the same initial spatial resolution with multiple refinement levels.

point calculation with time in the computation of geometrical information in AMR procedures. Nevertheless, the mass variation in each time step is at machine precision.

Summing up, our adaptive transport scheme is capable of accurately handling the divergent flow on adaptive meshes. The numerical error is nearly the same on non-adaptive meshes as on adaptive meshes and the scheme conserves mass in each time step. The heuristic gradient-based refinement criterion controls the mesh distribution by capturing the relevant tracer field and improves the efficiency of the numerical simulation. Better error estimators may further improve computational efficiency. The test case demonstrates that our adaptive transport scheme is able to be used in realistic simulations.

3.2.3 Non-Divergent Flow with Global Tracer Distribution: The Moving Vortices Test Case

The moving vortices test case is a challenging test case for AMR. Numerical accuracy on adaptive meshes and globally refined meshes is similar regardless of the feature of the flow when we use local tracer distributions as shown in Section 3.2.1 and 3.2.2. The moving vortices test case utilizes a global tracer distribution. To avoid global refinement in our AMR runs, the goal of our AMR scheme is to improve the local repre-

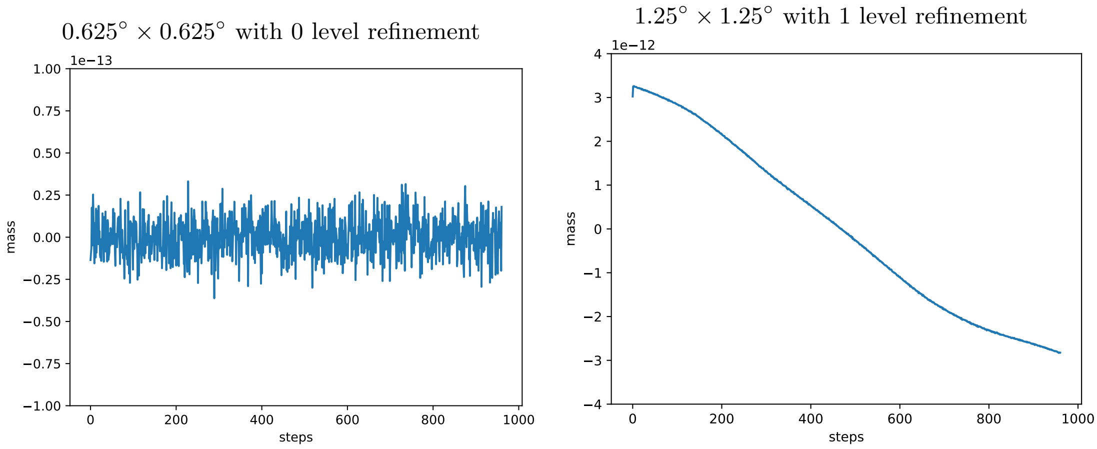

Figure 12. Evolution of mass on both non-adaptive meshes and adaptive meshes. The loss of mass arises from the accumulated floating point rounding error with time. The mass variation in each time step is at machine precision.

sentation of the tracer distribution in vortices instead of improving the numerical accuracy globally.

As the vortices in this test case develop with time, local refinement is not present at initial time steps. Our numerical experiments use low-resolution initial condition, which is different from experiments in Section 3.2.1 and 3.2.2.

The moving vortices test case allows us to mimic the setting in our targeted applications in ECHAM6. Our integrated adaptive transport scheme uses information from non-adaptive low-resolution dynamical core and parameterizations. Further, as the momentum equations are still solved on coarse resolutions by the spectral dynamical core, our AMR scheme needs to interpolate the wind field from the coarse mesh to the AMR mesh.

We use a coarse input wind field and a coarse initial tracer distribution in the test case. The wind field on higher resolution grids then has to be obtained by interpolation. To prevent numerical oscillations and maintain monotonicity, we use first-order bi-linear interpolation. To avoid excessive refinement and problematic interpolation around poles, we decide not to refine cells around the poles. We also expect such treatment leads to better efficiency.

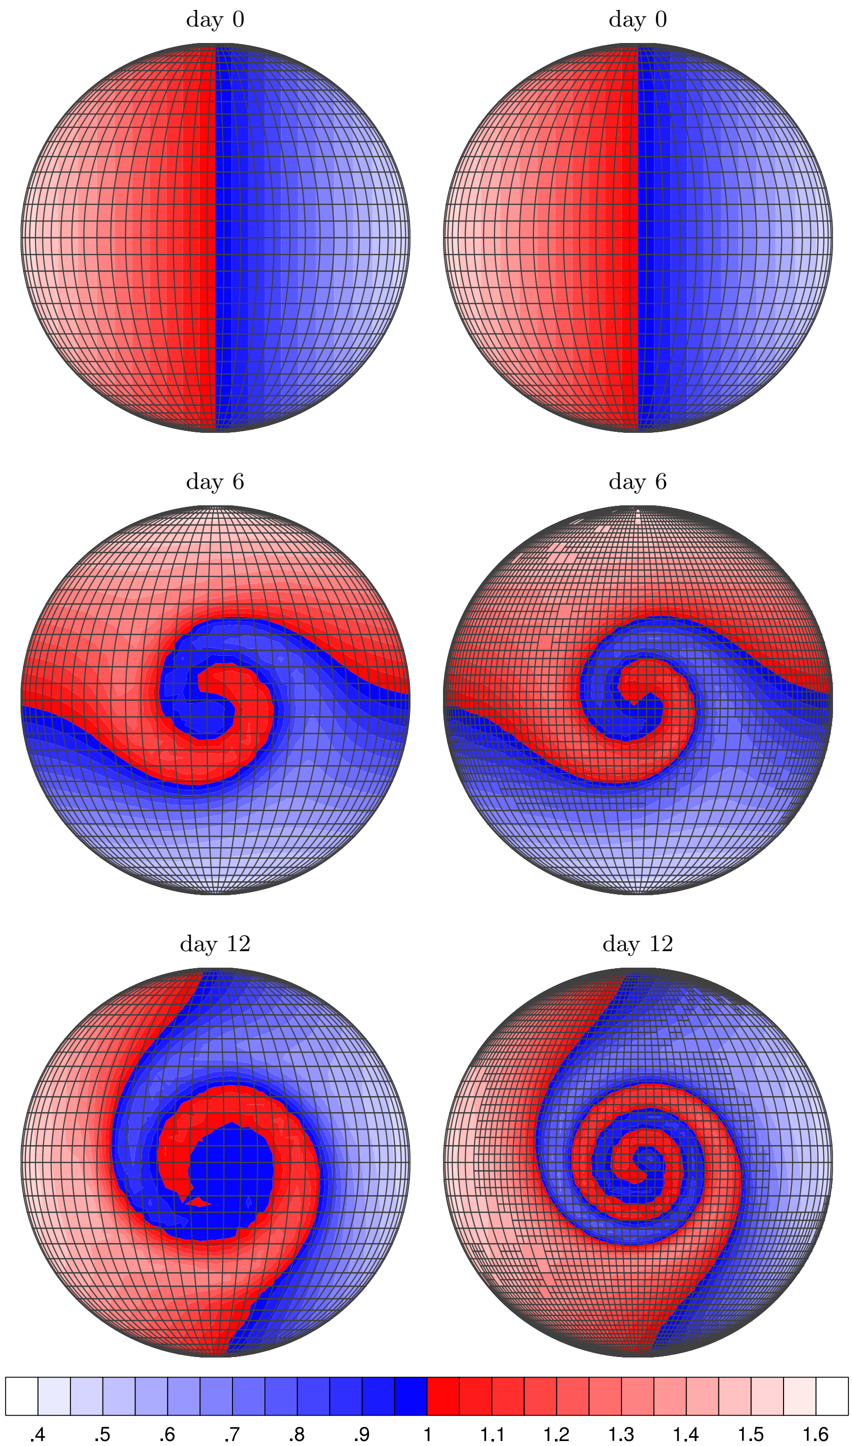

Figure 13. Numerical results of the moving vortices test case. The left panel shows the numerical results on a resolution of $5^\circ \times 5^\circ$ coarse grid. The right panel shows the numerical results on a resolution of $5^\circ \times 5^\circ$ coarse grid with one level refinement with interpolated wind field.

Compared to non-adaptive experiments, our AMR experiments lead to two sources of error: the error from coarse initial conditions and the error from wind interpolations. To investigate these errors, we examine three different settings. 1) We set up regular numerical experiments, where the initial condition and wind field is defined analytically on grid cells. 2) We run AMR experiments with one level and two level adaptive refinement as in previous sections, where coarse initial condition and interpolated wind field from initial refinement levels are used. 3) We also set up experiments using uniform refinement with coarse initial condition and wind interpolation. Note that a study on the sensitivity of wind interpolation on tracer fields were also performed in Behrens et al. (2000). We set the third experiment setting as reference solutions, which can show the errors arising from the coarse initial condition and wind interpolations. As a reference to the AMR experiments (setting 2), the coarse initial condition is consistent with the refinement levels in AMR experiments.

In all experiment settings, we always set $\alpha = \frac{\pi}{4}$ and test the numerical scheme with both large and small Courant numbers on various resolutions. On adaptive meshes, the refinement threshold for the gradient-based refinement criterion is $\theta_r = 0.8$ and the coarsening threshold is $\theta_c = 0.4$. The threshold in this test case is more relaxed than in the solid body rotation test case due to the strong deformation arising from the vortices.

We show snapshots of the numerical solution on $5^\circ \times 5^\circ$ coarse resolution and one level refinement in Figure 13. The refinement criterion captures the development of the vortices. Finer grids reduce the error around discontinuities induced by the vortices. The filaments of the tracer are not identifiable in low-resolution simulations but high-resolution simulations can capture the fine-scale feature in the tracer field such that we resolve finer filaments. Our adaptive transport scheme locally refines the regions where vortices appear. Our results indicate that AMR can improve local accuracy of numerical results even if the scheme can only access coarse grid information.

As shown in Figure 14, errors from the initial condition and wind interpolation contribute to the overall errors. Our results from the AMR setting and the uniform refinement setting show similar accuracy. However, regular numerical experiments without errors from the initial condition and wind interpolation show better results. A higher level of refinement with the same maximum spatial resolution on the mesh indicates a coarser

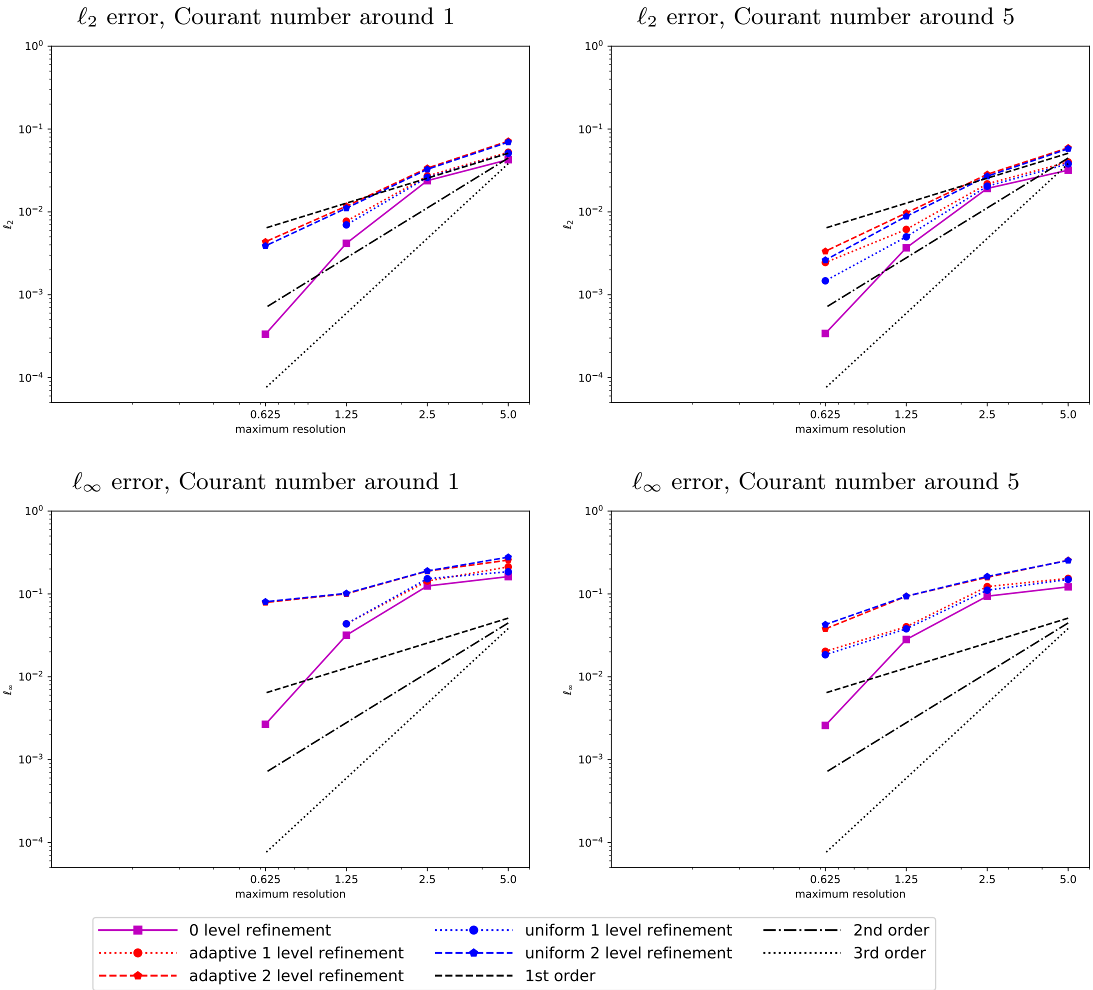

Figure 14. Convergence rate of the numerical results in the moving vortices test case on adaptive meshes.

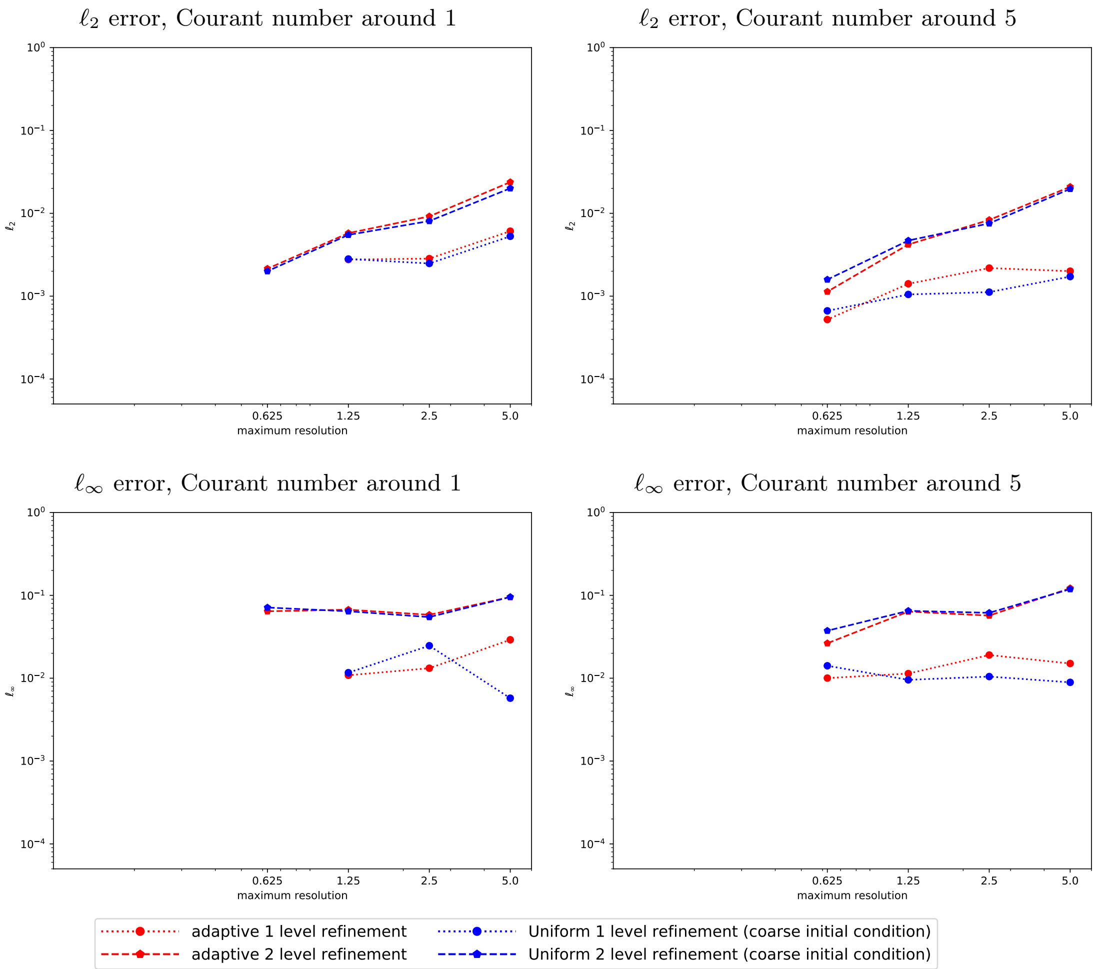

Figure 15. Differences of numerical errors between exact wind field and interpolated wind field in the moving vortices test case on adaptive meshes.

grid in the initial condition and wind field and correspondingly greater error than a lower level of refinement. We can also observe improved accuracy due to AMR compared to numerical results from low-resolution experiments.

To separate the influence of initial condition from that of the wind interpolation, we also conduct experiments using the exact wind field and coarse initial condition. We show the difference of numerical errors between the exact wind field and the interpolated wind field in Figure 15. We demonstrate that the inexact wind field leads to greater error than the coarse initial conditions.

Although the coarse initial distribution reduces the effect of refinement, using the high-resolution mesh still results in better numerical accuracy than only using the low-resolution mesh. Coarse input wind reduces the numerical accuracy. However, we still observe convergent and accurate numerical results using the AMR scheme. Our AMR scheme can improve the numerical accuracy using fewer grid cells than uniformly refined mesh when we integrate it into the tracer transport module into an existing model with coarse resolution simulations.

4 A Realistic Test Case: Simulation of Dust Transport

The tracer transport process exhibits multi-scale features in climate simulations. As indicated in Section 3, low-resolution simulations cannot represent fine-scale features of the tracer transport processes. Improving the local representation of the tracer transport scheme can, therefore, reduce at least one source of errors in climate models. On the other hand, the tracer transport process plays an important role in climate systems. The transported gases and aerosols have a significant impact on the state of climate through solar radiation (Carslaw et al., 2010). For example, carbon dioxide is one of the major driving factors of anthropogenic climate change. Volcanic ashes have a cooling effect on the global temperature. Hence, better tracer transport simulations can improve overall results in climate simulations.

We select dust to test our adaptive transport scheme in realistic settings. Dust has evident local origins like the Sahara desert and it can traverse across long distances while retaining local features as the atmospheric flow can lift dust to higher levels (Liu & Westphal, 2001). Emission and deposition parametrizations have less impact on higher level aerosols. Hence, dust simulations are suitable to demonstrate the advantages of using AMR.

We test our AMR scheme while maintaining a non-adaptive coarse climate model to which our AMR scheme is coupled in a one-way fashion. The one-way coupling prevents our tracer from interacting with other components of the climate model such that we can compare the difference between our adaptive tracer transport scheme and the original scheme using our conclusions from Section 3.

4.1 The Host Model: ECHAM-HAMMOZ

We integrate our adaptive tracer transport scheme into ECHAM6 without breaking the current code structure of ECHAM6. Further, the structure of ECHAM6 can also provide insight into numerical results of our simulation of dust transport. Hence, it is necessary to understand the model.

ECHAM6 is the atmospheric component of the earth system model, MPI-ESM (Stevens et al., 2013). It is composed of several components: the dynamical core, the physical parametrizations, and a land surface model, JSBACH.

The dynamical core solves hydrostatic primitive equations of the atmosphere, which describe the motion of air and assume absence of acceleration in the vertical. The dynamical core in ECHAM6 was originally derived from an early version of the atmospheric model developed at the European Center for Medium-Range Weather Forecast (Eliasen et al., 1970). ECHAM6 also applies a terrain-following coordinate to accommodate the varying orography at the bottom of the atmosphere. The terrain-following coordinate is a hybrid coordinate (Simmons & Burridge, 1981) leading to a non-orthogonal vertical mesh. Both the passive tracer transport scheme and the parametrizations in ECHAM6 are computed on a Gaussian grid using the flux-form semi-Lagrangian scheme, which we discussed in detail in Section 2. ECHAM6 also includes various parameterization schemes, including convection, cloud, radiation and vertical diffusion, etc. The land surface model comprises a class of parametrizations that provides the properties of land surface for other components of the climate model.

ECHAM-HAMMOZ is a coupled model that combines ECHAM6 and HAMMOZ since ECHAM6 has the flexibility to include various sub-models. HAMMOZ is a class of aerosol and atmospheric chemistry modules (Schultz et al., 2018) that predict the evolution of aerosols and trace gases. In our applications, we focus on the evolution of the dust concentration. ECHAM-HAMMOZ divides tracers into seven different modes (Vignati et al., 2004). These modes are dependent on the size, and solubility of the particles. There are four different modes for dust: Accumulation mode mixed (DU_AS), Coarse mode mixed (DU_CS), Accumulation mode insoluble (DU_AI) and Coarse mode insoluble (DU_CI). HAMMOZ describes the emission, diffusion, dry deposition, wet deposition, cloud scavenging and sedimentation of these tracers.

4.2 Tendency Equation of Dust Concentration

We replace the 2-D tracer transport scheme in ECHAM6 with our proposed AMR scheme. However, the evolution of dust concentration in a climate model is more complicated than a 2-D tracer transport equation. The large-scale temporal changes of dust concentration are not only controlled by tracer transport but also affected by various other parametrizations. The large-scale temporal changes of the tracer concentration are also referred to as the tendency of the tracer concentration.

In this section, we present the tendency equation of the dust concentration in ECHAM6. In addition, we also present our implementation when integrating our adaptive transport scheme to ECHAM6.

4.2.1 Numerical Treatment of Tendency Equation in ECHAM6

ECHAM6 describes the tendency equation of the tracer concentration using the following equation:

$$\frac{\partial \rho c}{\partial t} + \nabla \cdot (\rho c \mathbf{u}) = F. \tag{19}$$

Here ρ is the air density, c is the tracer concentration, the combination of ρc is the density of the tracer in the air, $\frac{\partial \rho c}{\partial t}$ is the tendency of the tracer concentration, $\nabla \cdot$ is the 3-D divergence operator, F represents external forcings. In climate models, the tracer concentration c represents the mixing ratio which is the mass of the aerosol or gas relative to the mass of dry air. The unit of mixing ratio is $\mathrm{kg\,kg^{-1}}$.

The forcing term includes the vertical diffusion, dust emission, dry deposition, wet deposition, sedimentation, and cloud scavenging process. The wet deposition process also involves the convective and cloud processes. Hence, the forcing term is a collection of parametrizations.

The tendency equation in the terrain-following hybrid coordinate is:

$$\frac{\partial p_s c}{\partial t} + \nabla \cdot (p_s c \mathbf{u}) = F \tag{20}$$

where p is the pressure and p_s is the surface pressure. Equation (20) also exhibits that the hybrid coordinate is based on the surface pressure. In hydrostatic systems, the hybrid coordinate prescribes a vertical pressure distribution.

The FFSL scheme in ECHAM6 leads to more diffusive results due to some modifications, which result in computationally less expensive scheme than the one presented in Section 3. For example, the FFSL scheme in ECHAM6 uses a first-order Godunov scheme as the inner operator and a third order piecewise parabolic method (PPM) as the outer operator instead of the third-order PPM for both inner and outer operators. The scheme includes limiters to ensure the positivity of the numerical results and averages over the longitude bands around the poles to avoid pole problems while we do not apply any limiters or special treatment around poles in Section 3.

4.2.2 Refinement Strategy

One of the benefits of integrating AMR into an existing model is that we do not need to implement and design a new model with the AMR technique. Rather, we can reuse most components of the existing model. In realistic dust simulations, we only need to replace the horizontal tracer transport scheme by our adaptive tracer transport scheme.

The hydrostatic primitive equations require the vertical integration of a column over each cell. Hence, for simplicity, instead of refining the mesh in 3-D, we only refine the horizontal 2-D mesh, obtaining locally smaller columns. 2-D refinement enables us to reuse the vertical tracer transport scheme without any modification.

As we integrate AMR into the passive tracer transport module without any modification in other components, the passive tracer transport module always gets wind, pressure and passive tracer concentration on a coarse grid. High-resolution wind can, therefore, only be obtained by interpolation from a coarse grid. Similar to the treatment of wind in Section 3, we use a bilinear interpolation. In the realistic test, we use the absolute value of ρc as a refinement criterion. When N tracers are simulated in ECHAM6, the refinement criterion is $\min(\rho c_i)$, where $i = 1, 2, 3, \ldots, N$. Hence, we take a refinement threshold of $1 \times 10^{-11}\,\mathrm{kg\,kg^{-1}}$ and a coarsening threshold of $1 \times 10^{-12}\,\mathrm{kg\,kg^{-1}}$.

4.3 Results of One-Way Coupling Dust Simulation

We test our adaptive tracer transport scheme with realistic dust concentration data using one-way coupling, i.e. we get coarse resolution wind and pressure as input data at each time step. During the simulations, we do not map the dust concentration back to the coarse resolution mesh used by other components. Therefore, the dust concentration does not affect other components of the climate model, especially pressure and wind field. Thus, we are able to investigate the effect of our AMR dust tracer transport using our conclusions from idealized tests in Section 3.

The dust concentration is always simulated on adaptive meshes. Since the parameterizations compute the tendency of tracer concentration in columns, our adaptive scheme can accommodate to use parameterizations.

4.3.1 Experiment Setting

In our one-way coupling experiments, parameterization schemes running on coarse resolution meshes should affect the dust concentration on adaptive meshes. Our implementation, refining columns, is aware of the original ECHAM6 parameterizations and is positivity preserving, leading to a compatible dust transport.

We can illustrate our treatment using a differential equation:

$$\frac{Dc_{\mathrm{AMR}}}{Dt} = F(X_{\mathrm{coarse}}, c_{\mathrm{AMR}}) \tag{21}$$

where $\frac{D}{Dt}$ is the material derivative, c_{AMR} is the tracer concentration of the AMR scheme, F is a parameterization scheme and X_{coarse} is a vector of variables involved in the parametrization scheme other than the tracer concentration. Therefore, our one-way coupling always uses coarse resolution parameters for parameterization schemes even if our tracer concentration is on higher resolution. We can achieve such implementation since parameterization schemes run within each column of the horizontal mesh. A flowchart in figure 16 illustrates this approach.

ECHAM6 provides a variety of options for the parameterization schemes. Although there are default settings for most parameterizations, we use some non-default options to simplify our experiment. In our experiment we use a vertical resolution of 31 layers, ($L31$). Hence, ECHAM6 does not compute the mid-atmosphere in our experiments.

In order to have the dust emission, we turn on the ECHAM-HAM submodel while we turn off the chemistry and MOZ submodel for simplicity. In our experiment, we also use the dust scheme proposed by Stier et al. (2005) and omit the additional Sahara and east Asia dust sources in the default settings.

We also set all agricultural, biogenic emission inactive, including forest fire and volcanic ashes. Hence, we only have emissions of dust species from the dust emission parameterizations. With this setting we simulate the dust evolution during the period of October 10 to October 31, 2006 as there are dust events at the Sahara desert during this month.

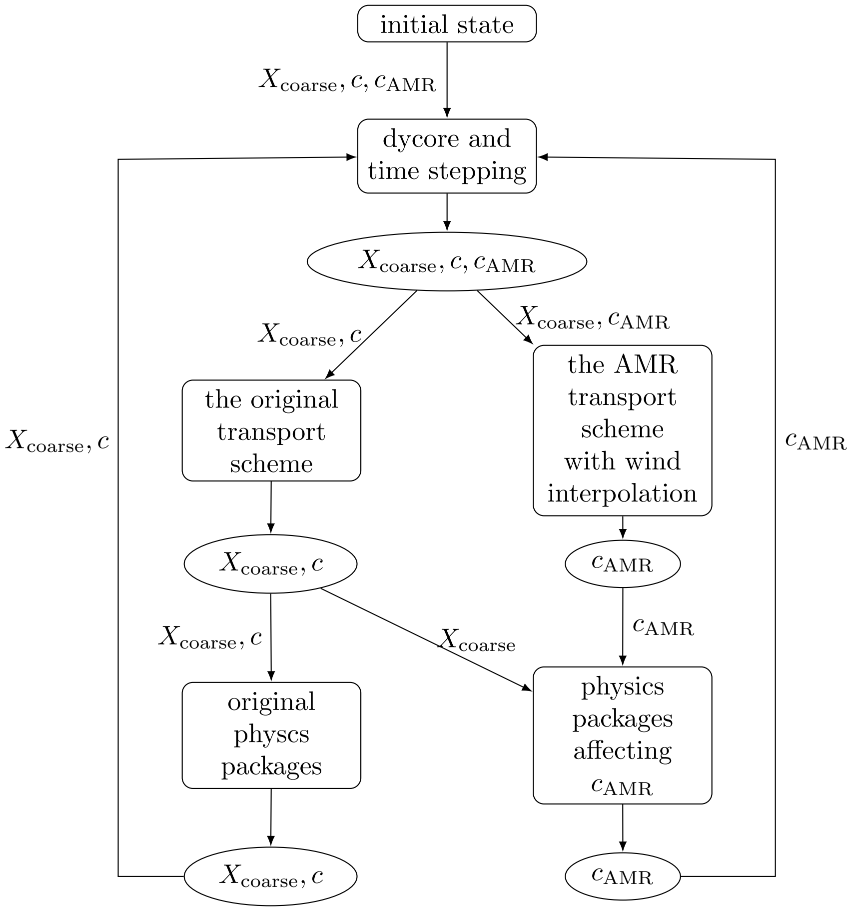

Figure 16. Illustration of our setting for one-way coupling experiment. c is the tracer concentration on the coarse resolution, X_{coarse} is a vector of variables other than the tracer concentration in the model on the coarse resolution, and c_{AMR} is the tracer concentration of the AMR scheme. The rectangle include modules/processes in the model, ellipse is the output of each module/process, arrows indicate the input variables in each module/process.

4.3.2 Comparison Between Low-Resolution and High-Resolution Simulations

We expect that high-resolution simulations can represent climate states with higher quality. High-resolution climate models not only represent the initial conditions better but also the boundary conditions, such as the topography and different types of land surface.

Our AMR scheme increases the resolution of the passive tracer transport scheme. However, our scheme can improve neither the initial condition nor the representation of the boundary conditions. Nevertheless, it is still of interest to compare the dust concentration on a low resolution of $T31L31$ and a higher resolution of $T63L31$ such that we can understand the difference between high-resolution simulations and low-resolution simulations.

We present the dust concentration of DU_AI in Figure 17. The Saharan air layer assumes large-scale systems can lift and transport dust up to a height of 5 km (Rodríguez et al., 2011). In order to capture the transport of dust without interference from the emission in lower levels, we show the dust concentration of DU_AI at 800 hPa both on $T31L31$ and $T63L31$ resolution.

Our results show that dust appears on 800 hPa after 3. Oct on coarse resolution simulations. The wind field transports dust westward toward the Atlantic ocean. After day 9, the dust concentration increases in East Asia and gradually moves south-westward.

However, the high-resolution simulation shows very different results. There is a high dust concentration at the east and west of North Africa respectively on 6. Oct while we cannot observe such high dust concentrations at low-resolution simulations. Although both dust simulations show a westward transport, the pattern of the dust distribution differs significantly. For example, hardly any dust disperses in east Asia in high-resolution simulations.

These significant differences arise from the parameterization and the resulting meteorological state. The emission of dust relies heavily on wind speed. It is difficult to attribute the difference to a single source due to the complexity of the interactions between various processes in the climate model.

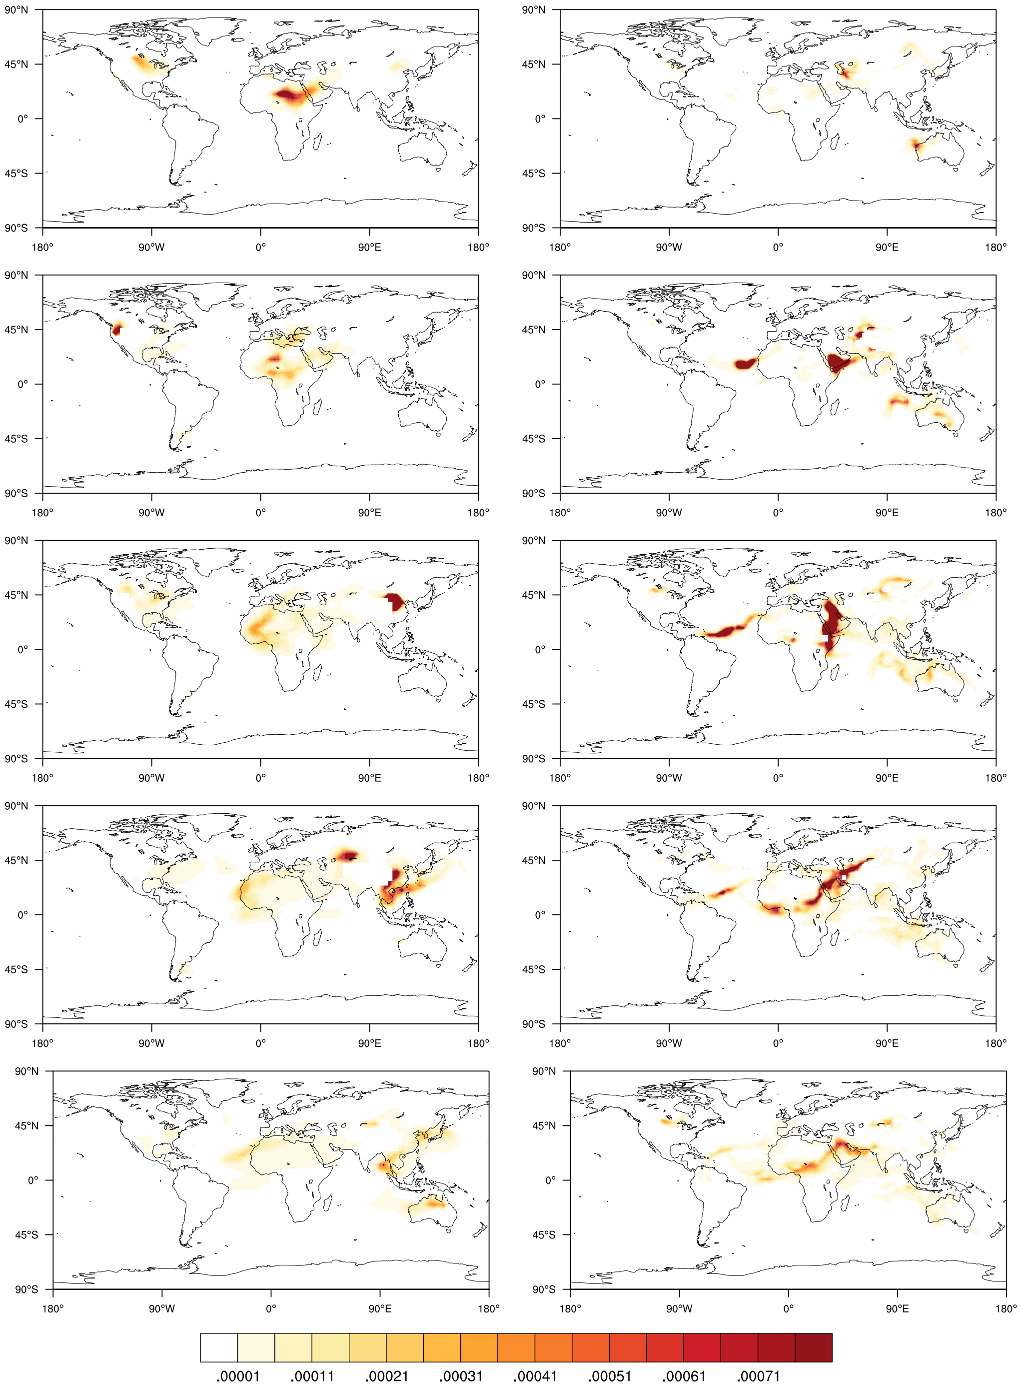

Figure 17. Dust concentration of DU_AI (mg kg^{-1}) at 800 hPa on 3rd, 6th, 9th, 12th and 15th October using model resolutions of $T31L31$ (left) and $T63L31$ (right). The dust concentration is masked due to high altitude in areas including the Tibet Plateau etc.

Nevertheless, our results demonstrate that low-resolution simulations with low-resolution initial conditions cannot compete with the results from high-resolution simulations with high-resolution initial condition and boundary conditions.

4.3.3 Comparison Between Low-Resolution and Adaptive Meshes

There are multiple sources for errors in low-resolution simulations. The coarse initial condition and boundary condition can lead to less accurate results while the coarse resolution dynamical core and parameterizations cannot resolve the finer features of the atmosphere. As discussed in Section 3, an interpolated wind field with coarse resolution initial condition can still improve the numerical accuracy of passive tracer transport schemes.

It is still promising that we can reduce one source of error from the coarse resolution climate simulations. Hence, we do not expect our results to be similar to results in high-resolution simulations. Although we do not have a reference solution for climate simulations on a coarse resolution, we can uniformly refine the mesh for passive tracer transport and interpolate wind to high resolutions.

We show results on uniformly refined meshes at the same period as the previous section in Figure 18. Compared to low-resolution simulations, we observe that uniformly refined meshes show less diffusive results. Dust concentration is higher than in low-resolution simulations while the filaments of the dust distribution are more obvious. It indicates that with reasonable refinement criteria, high-resolution dust transport is better than low-resolution dust transport. The parametrizations and the coarse dynamical core do not severely reduce the effect of high-resolution dust transport.

We take the uniformly refined mesh as the benchmark for our adaptive mesh refinement. Our results in Figure 19 show that AMR captures the appearance of dust and shows similar results on uniformly refined meshes and adaptive meshes, which demonstrate AMR can improve the accuracy of dust transport in the realistic test.

Our results show that integrating AMR into a passive tracer transport scheme can effectively reduce errors even if we do not use high resolution for other components. In order to capture the appearance of dust concentration, we use a strict refinement threshold, which leads to very large refinement regions. It is reasonable to choose a stricter refinement threshold or more appropriate refinement criteria in future applications.

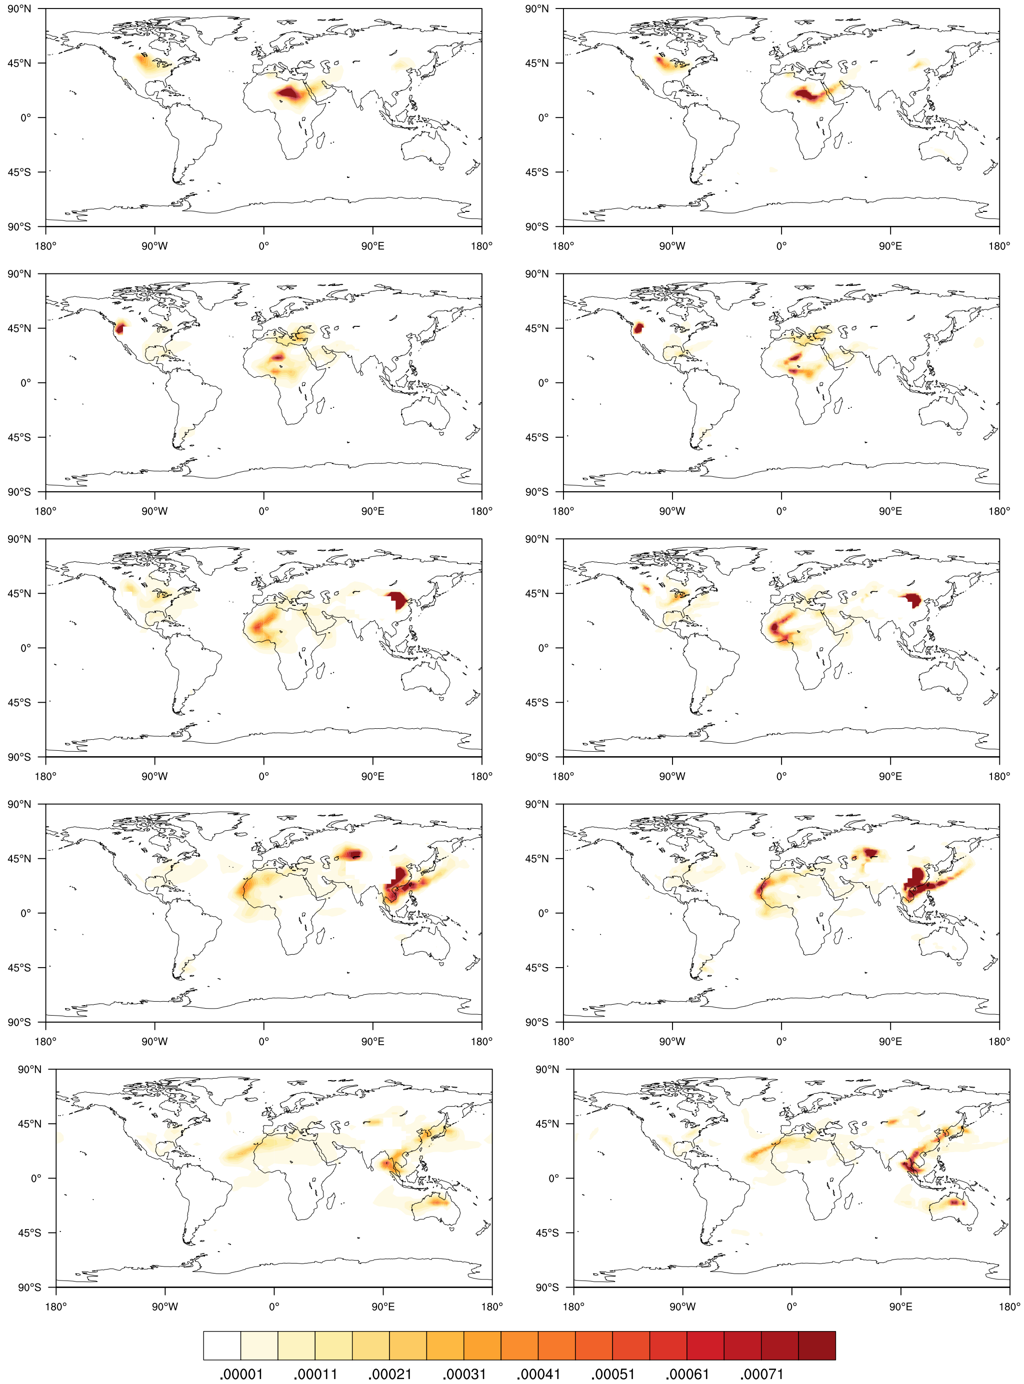

Figure 18. Dust concentration of DU_AI ($\mathrm{mg\,kg^{-1}}$) at 800 hPa on 3rd, 6th, 9th, 12th and 15th October on a model resolution of $T31L31$ using our modified transport scheme. The entire model runs on $T31L31$ on the left panel while the dust transport module has double resolution while the rest of the model is on $T31L31$ on the right panel.

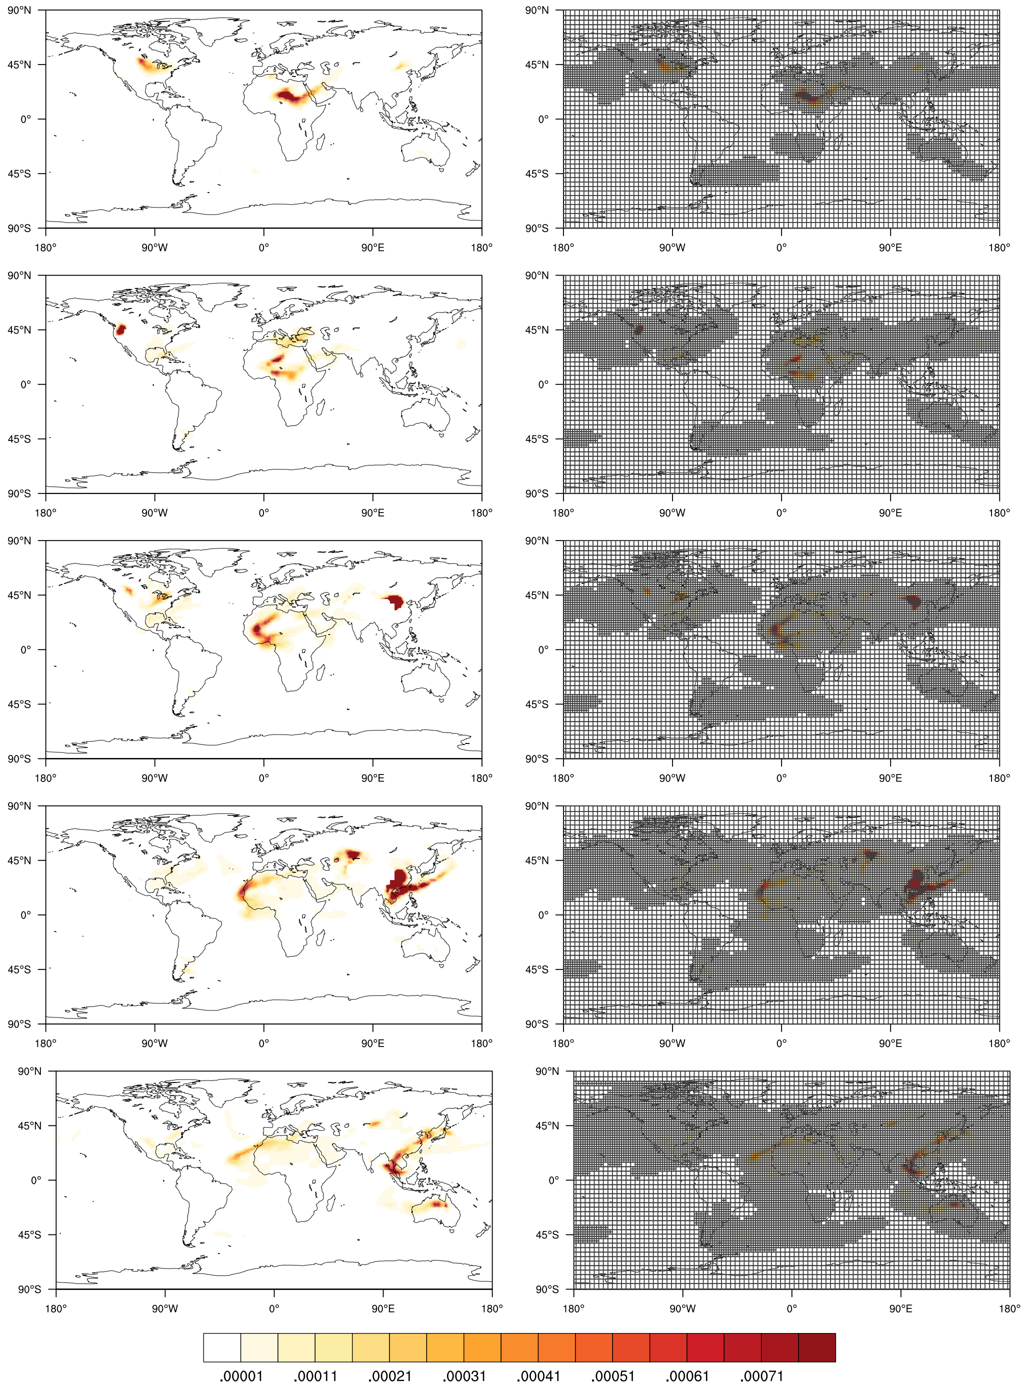

Figure 19. Dust concentration of DU_AI ($\mathrm{mg\,kg^{-1}}$) at 800 hPa on 3rd, 6th, 9th, 12th and 15th October on a model resolution of $T31L31$ using our modified transport scheme. The entire model runs on $T31L31$ on the left panel while the dust transport is on adaptive meshes and the rest of the model is on $T31L31$ on the right panel.

5 Conclusion

We propose a new approach towards adaptivity in climate models. Our method is different from the traditional AMR approach, which constructs a completely new climate model using AMR. Our approach overcomes the difficulty in integrating AMR into operational climate models. We integrate an AMR passive tracer transport scheme into an existing atmospheric model, ECHAM6. Improvements in a single component of a climate model should also improve the overall performance of the climate model. Partially integrating AMR into the existing climate model brings about an immediate improvement in accuracy and efficiency in operational climate simulations.

We use our approach to simulate dust transport processes in ECHAM6. Our AMR approach avoids mesh refinement of the entire globe and successfully captures regions where high-resolution meshes are necessary. High-resolution simulations improve the accuracy of dust transport processes but the general accuracy of the climate simulation is limited by the spatial resolution of other components.

Our idealized tests indicate that our AMR approach can potentially be as accurate as global high-resolution simulations when the tracer is present at local areas and the AMR scheme can access the exact wind field. Reducing local numerical errors can improve the overall accuracy of numerical solutions. Our AMR scheme leads to superior accuracy and efficiency compared to non-adaptive schemes.

Enabling AMR into existing climate models in each component relies on our investigation of several techniques: transport schemes, AMR strategies, and data structures, which is proposed by Chen et al. (2018). These techniques can be applied in a wider context than our applications.

Our modification to the widely used flux-form semi-Lagrangian (FFSL) transport scheme in ECHAM6 allows the transport scheme to be used on adaptive meshes while the transport scheme retains its important properties: dimensionally split, mass conservation, and semi-Lagrangian time stepping. Preserving the dimensionally split property results in the possibility of a fair comparison between the AMR scheme and the original scheme. Mass conservation is essential for climate models as it is unphysical to observe mass variation in transport processes. The semi-Lagrangian time stepping is particularly useful for AMR because the property can use a uniform time step on adaptive meshes without any stability issues. Hence, similar to the original FFSL scheme, our AMR scheme is a candidate for more complex systems (Lin, 2004; Jablonowski et al., 2009).

We also demonstrate the effectiveness of our proposed AMR strategy for dimensionally split schemes. Our AMR strategy ensures that high-resolution information is always transported on a high-resolution mesh, which guarantees the accuracy of numerical results. Thus, our AMR strategy results in accurate simulations as discussed in Section 3. Our modified FFSL scheme and AMR strategy lay a foundation for integrating AMR into existing models.

We expect that our results from dust simulations are applicable to other aerosols and gases as well. However, more rigorous investigations are needed. It is still of interest to explore a two-way coupling, where aerosols on adaptive meshes have an impact on processes such as cloud formation, radiation, and pressure, etc. The investigation on two-way coupling implies that we need to retain high-resolution information when we pass the information to low-resolution mesh. Averaging can lead to the loss of some fine-scale features. We require more sophisticated multi-scale methods to upscale high-resolution information to low-resolution meshes. These upscaling methods are a reverse of AMR, which upscales high-resolution information to low-resolution meshes (Simon & Behrens, 2018).

Our method may also be extended to more components of climate models. In addition, the implementation of our AMR schemes demands significant work on code optimizations and parallelizations for efficient operational climate models. Another possible use of AMR could be the dynamical coarsening of the mesh for a single component. Dynamical coarsening can circumvent the limitation of coarse initial conditions and parameterizations. However, it may require better data structures for it.

Our approach enables AMR component-wise in existing climate models, which reduces significant time of development compared to constructing a complete new AMR climate model.

Acknowledgments

This work was supported by German Federal Ministry of Education and Research (BMBF) as Research for Sustainability initiative (FONA); *www.fona.de* through Palmod project (FKZ: 01LP1513A). We also acknowledge support by the Cluster of Excellence CliSAP

(EXC177), Universität Hamburg, and Germany's Excellence Strategy – EXC 2037 'CLICCS – Climate, Climatic Change, and Society' – Project Number: 390683824, contribution to the Center for Earth System Research and Sustainability (CEN) of Universität Hamburg, both funded by the German Science Foundation (DFG). Besides, this work is also partially supported by the completion scholarship at Universität Hamburg. Data used for realistic dust simulations are obtained from DKRZ server and detailed description is available at https://redmine.hammoz.ethz.ch/projects/hammoz/wiki/V0002.

References

Behrens, J. (1996). An adaptive semi-lagrangian advection scheme and its parallelization. *Monthly weather review*, *124* (10), 2386–2395.

Behrens, J., Dethloff, K., Hiller, W., & Rinke, A. (2000). Evolution of small-scale filaments in an adaptive advection model for idealized tracer transport. *Mon. Wea. Rev.*, *128* , 2976–2982.

Berthet, S., Sfrian, R., Bricaud, C., Chevallier, M., Voldoire, A., & Eth, C. (2019). Evaluation of an online grid-coarsening algorithm in a global eddy-admitting ocean biogeochemical model. *Journal of Advances in Modeling Earth Systems*, *11* (6), 1759-1783. Retrieved from `https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2019MS001644` doi: 10.1029/2019MS001644

Carslaw, K., Boucher, O., Spracklen, D., Mann, G., Rae, J., Woodward, S., & Kulmala, M. (2010). A review of natural aerosol interactions and feedbacks within the earth system. *Atmospheric Chemistry and Physics*, *10* (4), 1701–1737.

Chen, Y., Simon, K., & Behrens, J. (2018). Enabling adaptive mesh refinement for single components in echam6. In *International conference on computational science* (pp. 56–68).

Colella, P., & Woodward, P. R. (1984). The piecewise parabolic method (PPM) for gas-dynamical simulations. *Journal of computational physics*, *54* , 174–201. doi: 10.1016/0021-9991(84)90143-8

Eliasen, E., Machenhauer, B., & Rasmussen, E. (1970). *On a numerical method for integration of the hydrodynamical equations with a spectral representation of the horizontal fields.* Kobenhavns Universitet, Institut for Teoretisk Meteorologi.

Iske, A., & Käser, M. (2004). Conservative semi-lagrangian advection on adaptive unstructured meshes. *Numerical Methods for Partial Differential Equations: An International Journal*, *20*(3), 388–411.

Jablonowski, C., Herzog, M., Penner, J. E., Oehmke, R. C., Stout, Q. F., Van Leer, B., & Powell, K. G. (2006). Block-structured adaptive grids on the sphere: Advection experiments. *Monthly weather review*, *134*(12), 3691–3713.

Jablonowski, C., Oehmke, R. C., & Stout, Q. F. (2009). Block-structured adaptive meshes and reduced grids for atmospheric general circulation models. *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, *367*(1907), 4497–4522.

Kessler, M. (1999). Development and analysis of an adaptive transport scheme. *Atmospheric environment*, *33*(15), 2347–2360.

Kopera, M. A., & Giraldo, F. X. (2015). Mass conservation of the unified continuous and discontinuous element-based galerkin methods on dynamically adaptive grids with application to atmospheric simulations. *Journal of Computational Physics*, *297*, 90–103.

Lauritzen, P. H. (2007). A Stability Analysis of Finite-Volume Advection Schemes Permitting Long Time Steps. *Mon. Wea. Rev.*, *135*, 2658–2673. doi: 10.1175/MWR3425.1

Lauritzen, P. H., Nair, R. D., & Ullrich, P. A. (2010). A conservative semi-lagrangian multi-tracer transport scheme (cslam) on the cubed-sphere grid. *Journal of Computational Physics*, *229*(5), 1401–1424.

Leonard, B., Lock, A., & Macvean, M. (1995). The nirvana scheme applied to one-dimensional advection. *Int. J. Numer. Methods Heat Fluid Flow*, *5*, 341–377. doi: 10.1108/EUM0000000004120

Leonard, B., Lock, A., & MacVean, M. (1996). Conservative Explicit Unrestricted-Time-Step Multidimensional Constancy-Preserving Advection Schemes. *Mon. Wea. Rev.*, *124*, 2588–2606. doi: 10.1175/1520-0493(1996)124⟨2588:CEUTSM⟩2.0.CO;2

Lin, S.-J. (2004). A vertically lagrangian finite-volume dynamical core for global models. *Monthly Weather Review*, *132*(10), 2293-2307. doi: 10.1175/1520-0493(2004)132⟨2293:AVLFDC⟩2.0.CO;2

Lin, S.-J., & Rood, R. B. (1996). Multidimensional Flux-Form Semi-Lagrangian Transport Schemes. *Mon. Weather Rev.*, *124*, 2046–2070. doi: 10.1175/1520

-0493(1996)124⟨2046:MFFSLT⟩2.0.CO;2

Liu, M., & Westphal, D. L. (2001). A study of the sensitivity of simulated mineral dust production to model resolution. *Journal of Geophysical Research: Atmospheres*, *106*(D16), 18099–18112.

Nair, R. D., & Jablonowski, C. (2008). Moving vortices on the sphere: A test case for horizontal advection problems. *Monthly Weather Review*, *136*(2), 699–711.

Nair, R. D., & Lauritzen, P. H. (2010). A class of deformational flow test cases for linear transport problems on the sphere. *Journal of computational physics*, *229*, 8868–8887. doi: 10.1016/j.jcp.2010.08.014

Nair, R. D., & Machenhauer, B. (2002). The Mass-Conservative Cell-Integrated Semi-Lagrangian Advection Scheme on the Sphere. *Mon. Wea. Rev.*, *130*, 649–667. doi: 10.1175/1520-0493(2002)130⟨0649:TMCCIS⟩2.0.CO;2

Rodríguez, S., Alastuey, A., Alonso-Pérez, S., Querol, X., Cuevas, E., Abreu-Afonso, J., ... De la Rosa, J. (2011). Transport of desert dust mixed with north african industrial pollutants in the subtropical saharan air layer. *Atmospheric Chemistry & Physics*, *11*(13).

Schultz, M. G., Stadtler, S., Schröder, S., Taraborrelli, D., Franco, B., Krefting, J., ... Wespes, C. (2018). The chemistry–climate model echam6.3-ham2.3-moz1.0. *Geoscientific Model Development*, *11*(5), 1695–1723. Retrieved from `https://www.geosci-model-dev.net/11/1695/2018/` doi: 10.5194/gmd-11-1695-2018

Simmons, A. J., & Burridge, D. M. (1981). An energy and angular-momentum conserving vertical finite-difference scheme and hybrid vertical coordinates. *Monthly Weather Review*, *109*(4), 758–766.

Simon, K., & Behrens, J. (2018). Multiscale finite elements through advection-induced coordinates for transient advection-diffusion equations. *arXiv preprint arXiv:1802.07684*.

Skamarock, W. C., & Klemp, J. B. (1993). Adaptive grid refinement for two-dimensional and three-dimensional nonhydrostatic atmospheric flow. *Monthly Weather Review*, *121*(3), 788–804.

Stevens, B., Giorgetta, M., Esch, M., Mauritsen, T., Crueger, T., Rast, S., ... others (2013). Atmospheric component of the mpi-m earth system model: Echam6. *Journal of Advances in Modeling Earth Systems*, *5*(2), 146–172.

Stier, P., Feichter, J., Kinne, S., Kloster, S., Vignati, E., Wilson, J., ... others (2005). The aerosol-climate model echam5-ham. *Atmospheric Chemistry and Physics*, *5*(4), 1125–1156.

Vignati, E., Wilson, J., & Stier, P. (2004). M7: An efficient size-resolved aerosol microphysics module for large-scale aerosol transport models. *Journal of Geophysical Research: Atmospheres*, *109*(D22).

Weller, H., Ringler, T., Piggott, M., & Wood, N. (2010). Challenges facing adaptive mesh modeling of the atmosphere and ocean. *Bulletin of the American Meteorological Society*, *91*(1), 105–108.

Williamson, D. L., Drake, J. B., Hack, J. J., Jakob, R., & Swarztrauber, P. N. (1992). A standard test set for numerical approximations to the shallow water equations in spherical geometry. *Journal of Computational Physics*, *102*(1), 211–224.