



Two conservative multi-tracer efficient semi-Lagrangian schemes for multiple processor systems integrated in a spectral element (climate) dynamical core

Christoph Erath^{1*}, Mark A. Taylor², Ramachandran D. Nair³

¹TU Darmstadt, Department of Mathematics, Dolivostraße 15, 64293 Darmstadt, Germany

²Sandia National Laboratories, Albuquerque, NM, 87185, USA

³National Center for Atmospheric Research, 1850 Table Mesa Drive, Boulder, CO 80305, USA

*Email address for correspondence: erath@mathematik.tu-darmstadt.de

Communicated by Luca Bonaventura

Received on June 12, 2015. Accepted on July 29, 2015.

Abstract

In today's atmospheric numerical modeling, scalable and highly accurate numerical schemes are of particular interest. To address these issues Galerkin schemes, such as the spectral element method, have received more attention in the last decade. They also provide other state-of-the-art capabilities such as improved conservation. However, the tracer transport of hundreds of tracers, e.g., in the chemistry version of the Community Atmosphere Model, is still a performance bottleneck. Therefore, we consider two conservative semi-Lagrangian schemes. Both are designed to be multi-tracer efficient, third order accurate, and allow significantly longer time steps than explicit Eulerian formulations. We address the difficulties arising on the cubed-sphere projection and on parallel computers and show the high scalability of our approach. Additionally, we use the two schemes for the transport of passive tracers in a dynamical core and compare our results with a current spectral element tracer transport advection used by the High-Order Method Modeling Environment.

Keywords: transport scheme, spherical geometry, cubed-sphere grid, conservative semi-Lagrangian, spectral element method, error, parallel scalability, performance

AMS subject classification: 65M08, 76M12, 65Z05, 68U99.

1. Introduction

In this work we give a review of two semi-Lagrangian schemes and present new results on the use of the two schemes for passive tracer transport in a state of the art dynamical core, i.e., we provide an efficient de-

parture grid algorithm and discuss efficiency on multiple processor systems and coupling possibilities.

In recent years, massively parallel petascale computers with hundreds of thousands of processor cores have become available. The next generation could have millions which allow for evermore precise modeling. However, only very scalable applications can benefit from these powerful machines.

In the last decades global spectral methods and finite volume methods have been the dominant methods used by atmospheric *dynamical cores* [1]. The dynamical core describes the fluid dynamical aspects of the atmosphere with the aid of partial differential equations. It is not surprising that finite volume schemes play an important role in fluid dynamical applications, since their formulation relies on the obligatory conservation of mass constraint. However, scalability beyond 10000 processors is a non trivial task for dynamical cores in atmospheric modeling [2].

In general, modeling on the sphere is not straightforward. The natural spherical coordinates, the latitude/longitude grids, have a singularity at the poles, i.e., clusters grid points at the poles, which also affects the performance of the code. Scalability, high-order accuracy, and the request to apply the schemes also on unstructured grids led to the adoption of Galerkin schemes in the atmospheric scientific community [3–5]. Since a few years the High-Order Method Modeling Environment (HOMME) [6], see also Section 4, more precisely its spectral element (SE) dynamical core [3], has been used by the Community Atmosphere Model (CAM, version 5.2 and higher) – CAM-SE. CAM-SE runs as a standalone atmospheric model but it is also the atmospheric component of the Community Earth System Model (CESM), a state-of-the-art climate model with atmosphere, ocean, land, and ice component models (see <http://www.cesm.ucar.edu>). To avoid the pole problem, HOMME uses the cubed-sphere geometry resulting from a gnomonic equiangular projection of the sphere [7]. This allows an efficient partition of the sphere on parallel platforms and leads to excellent scalability [2], i.e., tested for up to 170000 processors.

In today's atmospheric modeling transport becomes a dominating factor of the total computational costs; e.g., more than 100 tracers are used in the chemical version of CAM. Besides the performance question it is a non trivial task to provide a numerical scheme such that the solution satisfies all the criteria of the (chemistry) climate modeling community: conservation of mass (also over long simulation times), positivity (negative values are not acceptable), accuracy, efficiency, parallel scalability, preservation of functional relations. The default time discretization in HOMME for the tracer transport is based on an explicit Runge-Kutta approach [8]. However, this approach requires three communications per time step with a relative small

time step, both caused due to the explicit Runge-Kutta method. Especially for multi-tracer transport this Eulerian type (with a fixed mesh) method is computationally expensive and is considered to be a major computational bottleneck of the HOMME dynamical core in the chemistry community.

Therefore it is an ongoing work to find alternative transport schemes. Another numerical method class to realize the transport problem in fluid dynamics is of Lagrangian type, where one follows individual parcels along their trajectories of the fluid, i.e., the mesh travels in the fluid. A combination of the Eulerian and Lagrangian idea is called Arbitrary Lagrangian-Eulerian (ALE) scheme and tries to benefit from both approaches. Semi-Lagrangian schemes belong to this family. Roughly speaking they do an interpolation from a Lagrangian mesh to a regular Eulerian mesh at every time step. Note that a semi-Lagrangian method such as SLICE [9] is very efficient for global multi-tracer transport problems and is based on the conservative cascade approach [10], which is particularly designed for structured lat/lon spherical grid systems. However, the dimension-splitting cascade remapping algorithm is not suitable for cubed-sphere type patched domains. Therefore, HOMME considers other semi-Lagrangian type schemes. Currently, three are implemented, all of them at a different stage of development. A scheme based on the traditional philosophy of the semi-Lagrangian idea uses a spectral element reconstruction followed by a global optimization to preserve conservation of mass and monotonicity (with a divided and conquer strategy) [11]. In this work we give a review of the conservative SPectral-Element Lagrangian Transport (SPELT) scheme in HOMME [12] and the Conservative Semi-Lagrangian Multi-tracer transport scheme (CSLAM) in HOMME [13]. Both schemes in HOMME are mass conservative, multi-tracer efficient, third order accurate, scalable on parallel computers and have a positivity option. For more details see Section 2 and Section 3, respectively. Furthermore, we will present a new coupling approach to the spectral element method for the first time, e.g., using the semi-Lagrangian schemes for passive tracer transport whereas the rest of the dynamical core calculation is done by the original spectral element method. For this, we introduce a special adapted departure grid algorithm for the HOMME system.

This paper is organized as follows; Section 2 describes the flux-form based multi-tracer efficient and conservative spectral-element Lagrangian transport scheme (SPELT). Section 3 gives a brief overview of the conservative semi-Lagrangian multi-tracer transport scheme (CSLAM). In Section 4 we introduce the High-Order Method Modeling Environment (HOMME) and describe the SPELT and CSLAM algorithm in this environment. Numerical examples show for both schemes the third order accuracy, the multi-

tracer efficiency, and the high scalability on parallel computers. An economical departure grid algorithm for HOMME and the spectral-element grid is derived in Section 5. Finally, Section 6 shows first results using the semi-Lagrangian schemes for passive tracer transport with spectral-element dynamics. Some conclusions can be found in Section 7.

2. SPectral-Element Lagrangian Transport

The SPELT scheme described in [12] is based on the multi-moment semi-Lagrangian scheme of [14]. SPELT is a third order accurate flux form scheme, which uses a biquadratic polynomial reconstruction (2D) on a single Eulerian cell. This local stencil avoids the use of a wider halo region. Together with an easy search algorithm, which is needed to find the fluxes, the SPELT approach is very attractive for non-uniform grids. In [12] we show that the scheme works on non-uniform quadrilateral grids with a positivity preservation option. Additionally, we discuss how the scheme can be made monotonic with a flux-corrected transport approach in the spirit of [15]. The implementation in HOMME and some high performance aspects are discussed later in Section 4. First we want to describe the basic feature of the scheme on a uniform Cartesian grid.

The flux-form transport equation for a scalar $\psi(x, y, t)$ in 2D Cartesian (x, y) -plane, without a source or sink, reads

$$(1) \quad \frac{\partial \psi}{\partial t} + \frac{\partial(u\psi)}{\partial x} + \frac{\partial(v\psi)}{\partial y} \equiv \frac{\partial \psi}{\partial t} + \nabla \cdot \mathbf{F} = 0 \quad t \in (0, T],$$

where $(u, v) =: \mathbf{u}$ is the wind velocity vector, $\mathbf{F} = (\psi u, \psi v)$ is the flux, $(0, T]$ the time interval with $T > 0$, and the initial condition is prescribed as $\psi(x, y, t = 0) = \psi_0(x, y)$ at time $t = 0$. However, for non-divergent flow fields, a regular semi-Lagrangian approach, which is not constrained to be conservative, uses the following simple form:

$$(2) \quad \frac{D\psi}{Dt} = 0, \quad \Rightarrow \quad \psi^{n+1} = [\psi^n]^*,$$

where the second equation indicates a semi-Lagrangian time discretization, and $D/Dt := \partial_t + u \partial_x + v \partial_y$ is the Lagrangian (material) derivative. Here, ψ^{n+1} is the estimate of ψ at the new time-level t^{n+1} from known values ψ^n at time $t^n = n\Delta t$, where Δt is the time step. Usually ψ^{n+1} is computed by an interpolation at the foot of the trajectory (upstream position), in (2) this is denoted as $[\cdot]^*$. SPELT benefits from both formulations, i.e., (1) ensures mass conservation of a quantity in a cell and the classical semi-Lagrangian scheme (2) is used to get a high order scheme. Let us consider a domain

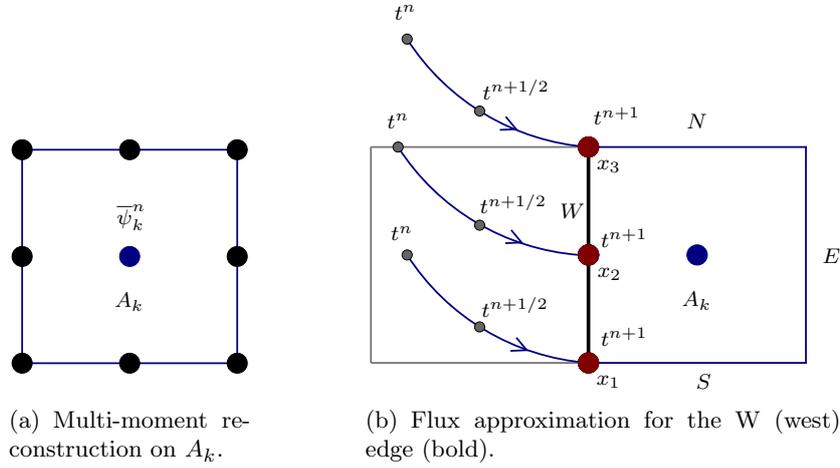


Figure 1. Figure (a) shows the multi-moment polynomial reconstruction stencil of SPELT, which has nine degrees of freedom on each cell. The quantity $\bar{\psi}_k^n$ is the prognostic variable. It is used for the reconstruction but updated by a classical flux form semi-Lagrangian approach to ensure mass conservation. The eight points on the edges are updated through a regular semi-Lagrangian approach along the trajectories. Figure (b) shows the flux approximation along the characteristics during the time $t \in [t^n, t^{n+1}]$ by the 3-point Simpson's rule. The flux on the west (W) edge (bold) of the cell A_k is also approximated with the aid of a 3-point Simpson's rule.

$\Omega \subset \mathbb{R}^2$. The triangulation of Ω into quadrilaterals is denoted by \mathcal{T} . If we integrate (1) spatially over the cell $A_k \in \mathcal{T}$ and temporally over the time interval $[t^n, t^{n+1}]$ we get

$$\int_{t^n}^{t^{n+1}} \int_{A_k} \left(\frac{\partial \psi}{\partial t} + \nabla \cdot \mathbf{F} \right) dx dy dt = 0.$$

If we write $\bar{\psi}_k^n$ for the integral mean of $\psi(x, y, t^n)$ in the cell A_k at time t^n the divergence theorem leads to the finite volume formulation

$$(3) \quad \bar{\psi}_k^{n+1} |A_k| = \bar{\psi}_k^n |A_k| - \int_{t^n}^{t^{n+1}} \left(\oint_{\partial A_k} \mathbf{F} \cdot \mathbf{n} ds \right) dt,$$

where \mathbf{n} is the unit normal vector on the boundaries ∂A_k of A_k pointing outward with respect to A_k . Note that (3) describes the mass balance, i.e., the mass at time t^{n+1} depends on the mass at time t^n plus some inflow/outflow fluxes. The accuracy and efficiency (stability) of the discretization of (3) depend on the line and time integrals. Following [14] we use a multi-moment reconstruction with nine degrees of freedom on each cell, see Figure 1(a). This polynomial reconstruction includes eight point values on the corners

and edges on each quadrilateral A_k , which are evolved in time and are updated through (2). The cell average $\bar{\psi}_k^n$ is the ninth degree of freedom. It is updated through (3) and is at the same time the conservative prognostic variable we are looking for. The explicit form of the coefficients in terms of nine values can be found in [16]. We note that the reconstruction is globally continuous. The last term in (3) can be written as

$$(4) \quad \int_{t^n}^{t^{n+1}} \left(\oint_{\partial A_k} \mathbf{F} \cdot \mathbf{n} ds \right) dt = \oint_{\partial A_k} \tilde{\mathbf{F}} \cdot \mathbf{n} ds,$$

where $\tilde{\mathbf{F}} \cdot \mathbf{n}$ may be interpreted as exact evolution of fluxes on the boundaries ∂A_k along the characteristics during the time $t \in [t^n, t^{n+1}]$ as indicated in Figure 1(b) for one edge (bold, west edge) of ∂A_k . We use a 3-point Simpson's rule to get

$$\tilde{\mathbf{F}} = \int_{t^n}^{t^{n+1}} \mathbf{F}(t) dt \approx \frac{\Delta t}{6} (\mathbf{F}(t^n) + 4\mathbf{F}(t^{n+1/2}) + \mathbf{F}(t^{n+1})).$$

Here, the fluxes at the departure point at t^n and the midpoint at $t^{n+1/2}$ are computed using the polynomial reconstruction. Since the reconstruction is locally defined on each cell we need to know the cells which contain these points, see Figure 1(b). Note that the departure and midpoint points are for all tracers the same. Therefore we have to do this search only once per time step, which makes the scheme multi-tracer efficient. For the final flux integral over one edge of $\tilde{\mathbf{F}} \cdot \mathbf{n}$ in (4) we use again the 3-point Simpson's rule, i.e., we need three fluxes $\tilde{\mathbf{F}}_{x_i}$, $i = \{1, 2, 3\}$ along the trajectories arriving in x_1 , x_2 , and x_3 , respectively, see Figure 1(b). Thus the flux at the west edge (W) reads

$$F_W = \frac{h_W}{6} (\tilde{\mathbf{F}}_{x_1} + 4\tilde{\mathbf{F}}_{x_2} + \tilde{\mathbf{F}}_{x_3}),$$

where h_W denotes the length of the edge W. Applying the same approximation for the other edges $\{S, E, N\}$ of A_k the final third order SPELT scheme reads

$$(5) \quad \bar{\psi}_k^{n+1} |A_k| = \bar{\psi}_k^n |A_k| - (F_W + F_S + F_E + F_N).$$

Note that the integral approximation can be done also with Gaussian quadrature points. The number of point searches and polynomial evaluations depend on the order of the quadrature rule. However, flux schemes have the strong property that they ensure mass conservation no matter how we approximate fluxes. Nevertheless, it is important to know for parallelization that if the flux approximation on an edge E is done separately

Scalable, conservative, multi-tracer efficient semi-Lagrangian schemes

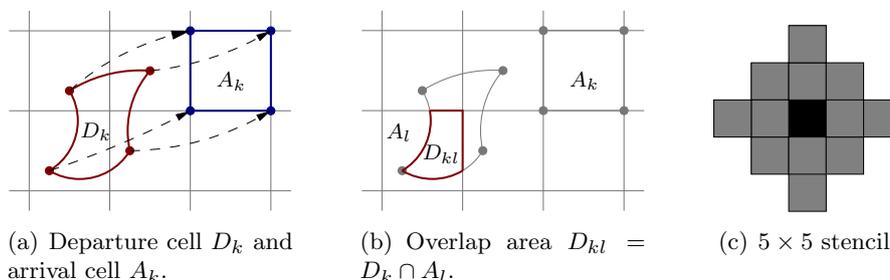


Figure 2. The departure cell D_k moves in one time step to the (Eulerian) arrival cell A_k in (a). For simplicity, we approximate the departure cell by connecting the departure points with straight lines. The overlap areas as in (b) are used to calculate the integral over the departure cell D_k . Figure (c) shows the 5×5 stencil, which is needed for the reconstruction in one cell (black) for the third order accurate CSLAM scheme.

on cells, which share E , it has to be the same up to machine precision. For more details on the flux approximation, in particular in the more difficult case of non-uniform grids, we refer to [12].

There, one can also find details on a quasi-monotone semi-Lagrangian (QMSL) filter and the flux-corrected transport (FCT) philosophy. The QMSL filter is a classical a posteriori filter, local, and very easy to implement. It guarantees that the newly computed point values are always within legitimate bounds and helps to control unphysical oscillations. However, QMSL is not enough to preserve positivity. The FCT technique was introduced by [15] to get monotonicity and a simple positivity preservation approach can be found in [17]. In this work we will apply the simple positivity preservation and the QMSL filter.

3. Conservative Semi-LAgrangian Multi-tracer transport scheme

In this section we give a brief overview of CSLAM [13,18]. CSLAM is an incremental remapping algorithm based on the idea of the works [19–21].

Note that (1) can be cast equivalently in the following Lagrangian form

$$(6) \quad \frac{D}{Dt} \int_{A(t)} \psi \, dx \, dy = 0,$$

where $A(t)$ is an arbitrary area (volume) moving with the fluid in which the fluid density ψ evolves in time along the Lagrangian trajectories (characteristics). CSLAM is based on the upstream semi-Lagrangian method. That means parcels that end up on the regular Eulerian grid \mathcal{T} are considered. Let us write $A_k \in \mathcal{T}$ for the arrival cell with the corresponding departure cell D_k , i.e., D_k moves in one time-step to A_k , see Figure 2(a).

We remark that we fix the arrival grid \mathcal{T} (Eulerian), thus our scheme is a semi-Lagrangian scheme with backward trajectories.

A temporal discretization of (6) along the characteristics and approximation of ψ in D_k leads to the discrete scheme

$$(7) \quad \overline{\psi}_k^{n+1} |A_k| = \int_{D_k} \psi_{k,rec}^n dx.$$

Here, $\overline{\psi}_k^{n+1}$ is the average tracer in A_k at time step t^{n+1} and $\psi_{k,rec}^n$ a certain reconstruction. To get a high order numerical scheme, $\psi_{k,rec}^n$ has to be a high order reconstruction from the cell averages of the previous time step t^n . CSLAM uses a \mathcal{T} -piecewise two-dimensional polynomial of degree two, where the coefficients are defined in terms of a Taylor expansion. The gradient and curvature are approximated using finite differences from the known cell averages $\overline{\psi}_k^n$. Note that we need this approximation in each coordinate direction as well as the cross derivative. The stencil, which is needed for third order accuracy, is shown in Figure 2(c). The constant term in the reconstruction is chosen such that the integral of the reconstruction function over A_k equals the mass in that cell. For more details we refer to [22]. Since this reconstruction is not globally continuous, one must split the integral over D_k in (7) into the sum of integrals over overlap areas between D_k and the regular Eulerian grid. Figure 2(b) shows one part of the integral over $D_{kl} = D_k \cap A_l$. For the CSLAM scheme the integral over D_{kl} is transformed into line integrals using the divergence theorem. Finally, this allows us to split the integral in (7) into weights $\omega_{kl}^{(i,j)}$, which are functions of the coordinates of the vertices of D_{kl} , and the reconstruction coefficients $c_l^{(i,j)}$. Then the third order CSLAM scheme reads

$$\overline{\psi}_k^{n+1} |A_k| = \sum_{l=1}^{L_k} \left[\sum_{i+j \leq 2} c_l^{(i,j)} \omega_{kl}^{(i,j)} \right].$$

The weights $\omega_{kl}^{(i,j)}$ can be re-used for each additional tracer which makes the scheme multi-tracer efficient, since an expensive search algorithm to define D_{kl} has to be done only once per time step. We remark, that conservation of mass relies not only on the constant term of the reconstruction $c_l^{(i,j)}$ but also on satisfying integral constraints. In particular, the CSLAM scheme relies on the analytical calculation of the interior line integrals $\partial D_{kl} \cap \partial A_l$. However, this analytical expressions are numerically unstable for high-resolution meshes on the sphere, at least for the third order scheme. The approximation of these integrals by a more robust Gaussian quadrature might lead to a stable scheme but also in losing of conservation of mass. The work

Scalable, conservative, multi-tracer efficient semi-Lagrangian schemes

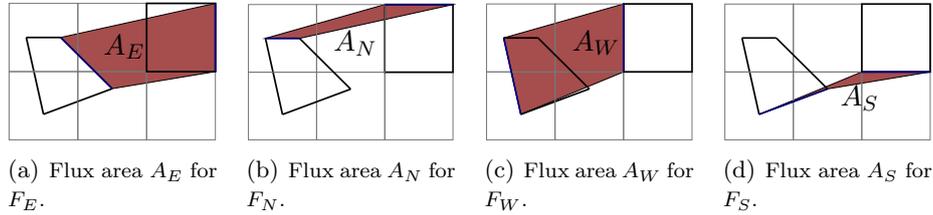


Figure 3. Areas for the fluxes in different directions for the flux-from CSLAM introduced in [25].

in [23] discusses this problem in more details for general high-order high-resolution rigorous remapping schemes on the sphere. Furthermore, that work provides a stability technique, referred as enforcement of consistency, which still preserves mass conservation and the multi-tracer efficiency. This general strategy for remapping schemes is confirmed on a test example with CSLAM.

As long as the upstream cells D_k are simply connected, the scheme is stable for long time steps ($CFL > 1$). To get monotonicity, one can apply the simple 2D reconstruction function filter from [24]. CSLAM scheme is extensively tested with various standard benchmark test cases of solid-body rotation and deformational flow, incl. on spherical geometry, see e.g. [18].

Remark 3.1. It is also possible to cast the above ideas to formulate a semi-Lagrangian scheme based on the flux form (1) transport equation. This has been discussed and tested in [25]. The scheme formulation is the same as for SPELT; see (5). But the fluxes F_X with $X = \{E, N, W, S\}$ are calculated with the same techniques as for the rigorous CSLAM scheme. The flux areas A_E, A_N, A_W, A_S are indicated in Figure 3. That means that the integral over A_X with $X = \{E, N, W, S\}$ can be split into weights and reconstruction coefficients to get a multi-tracer efficient flux form scheme. Although the flux version is much more expensive, i.e., one has to find four searching areas instead of only one, a flux form scheme has in particular some essential advantages, as was noted already for the SPELT scheme in Section 2.

4. Semi-Lagrangian schemes in HOMME

In this section we describe very briefly HOMME and give an overview of the implementation techniques from SPELT and CSLAM in HOMME. Unlike the global scheme SLICE [9], the conservative semi-Lagrangian schemes SPELT and CSLAM are local methods, which are well suited for the HOMME data structure, i.e., the cubed-sphere grid and the spectral

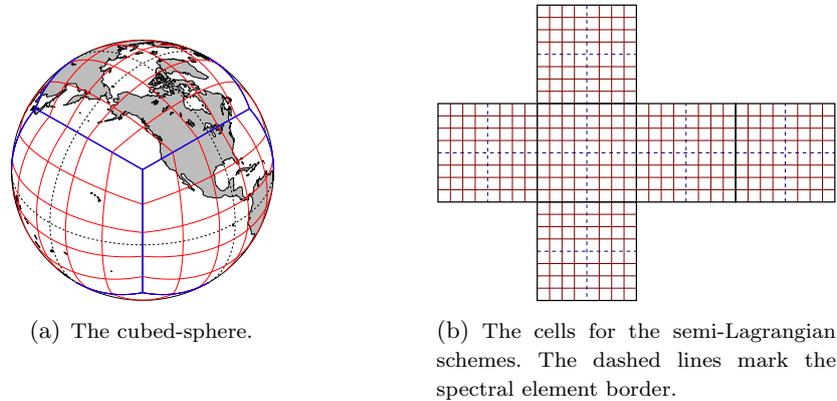


Figure 4. The cubed-sphere in (a) and the finite volume grid on the spectral element grid on the cube faces in (b), which is projected through an gnomonic projection onto the sphere. The cells are not equidistantly spaced as it is shown here for the sake of simplicity.

elements. Some error and scalability discussions are reported in the last subsection.

4.1. *The High-Order Method Modeling Environment*

The domain decomposition (horizontal) in HOMME follows along the cubed-sphere grid [7] to avoid the pole problem. That means, an inscribed cube is projected through an equiangular central (gnomonic) projection on the surface of the sphere, see Figure 4(a). For Galerkin schemes the six faces of the cube can be easily divided into finite-elements. They build the natural base for the partition on parallel platforms. Therefore, the parallel strategy in HOMME relies on the spectral elements, i.e., all cube faces are divided in quadrilateral elements. Note that this allows also locally refined meshes, which are in particular of interest, if one wants to resolve local phenomena such as singularities and shock waves. A standard configuration is based on an equidistant mesh on each face, which leads to an almost equidistant mesh on the sphere, see Figure 4(a). For each spectral element 4 or 5 Gaussian-Legendre-Lobatto (GLL) points are used. That means the number of unknowns for the spectral element scheme for each processor is at least 16 or 25 if we run the model with a minimum of one spectral element per processor. The work in [2] shows that this approach in HOMME is highly scalable, up to 170000 processors on the Oak Ridge LCF Cray XT5 JaguarPF system. Moreover, [2] gives a short summary on the most important components/numerical schemes used for the dynamical core for CAM-SE, including the vertical discretization method.

Scalable, conservative, multi-tracer efficient semi-Lagrangian schemes

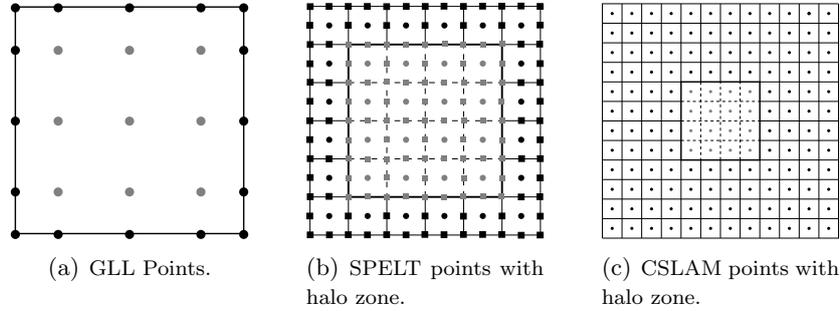


Figure 5. In Figure (a) we show 5×5 GLL points for a spectral element and the exchange points (black). Figure (b) presents the reconstruction points for SPELT on a 4×4 cell grid on the spectral element including the halo zone (black points) with depth one (cell), which allows a $CFL \leq 1$. And finally, Figure (c) shows the CSLAM points including the halo zone (black points) with depth four (cells), which allows a $CFL \leq 1$.

The default tracer transport scheme in HOMME is built on the native spectral element discretization (locally conservative) with an explicit Runge-Kutta based approach [8]. To avoid oscillations the scheme in the model comes with an optimization based monotone limiter for the hp spectral method. Additionally, it is coupled with a hyper-diffusion operator. However, this leads to a relative small time step constraint and three communications per time step. Therefore, for multi-tracer transport the current configuration is computationally expensive.

4.2. SPELT and CSLAM in HOMME

With a cubed-sphere grid the transfer of SPELT and CSLAM onto the sphere is straight forward but nevertheless not a trivial task. The search algorithm, the reconstruction and so on in gnomonic coordinates is much more complicated. We refer to the works [12–14,18]. If one wants to parallelize these algorithms, the situation is even more delicate, which we want to review in this subsection.

The algorithm of SPELT and CSLAM in HOMME are described and intensively discussed in [12] and [13], respectively. Both schemes require only one communication per time step with the concept of an extended element. SPELT is designed for non-uniform grids since it uses a point wise search and a local (one cell) reconstruction [12], i.e., the halo zone requires one cell in the extended element. In this work we consider a SPELT version on the finite volume grid of HOMME, i.e., each spectral element is divided into cells, see Figure 4(b). The CSLAM algorithm works on the same grid. We remark that the 5×5 reconstruction stencil is rather complicated to

implement on a cube, especially on the cube corner. Therefore, CSLAM requires a halo zone with four cells in the extended element, if one also wants to consider the cross terms in the reconstruction and avoid additional communication. Furthermore, the search process involved in the CSLAM algorithm is computationally intense and challenging in complex domains. The computational costs for the enforcement of consistency version, see Section 3 or [23], are slightly higher than for the original CSLAM but are negligible.

An extended element for SPELT and CSLAM can be seen in Figure 5(b) and Figure 5(c), respectively. Here, the filled circles are the cell averages $\bar{\psi}_k^n$ in the halo zone. The basic idea is to avoid additional communication, since we calculate certain quantities twice, e.g., for SPELT we calculate the fluxes on a spectral element edge once in the halo cell and once in the cell of the corresponding element on different processors. This can be done up to machine precision which is important to ensure mass conservation. We remark that this was never an issue in practice. However, the technique of an extended element increases the communication bandwidth. For CSLAM the CFL is restricted to $\text{CFL} \leq 1$ to ensure communication between neighboring spectral elements. For SPELT one could relax this restriction to $\text{CFL} \leq 3$ on the cost of a higher bandwidth and a higher order of quadrature rule [14]. However, we limit the $\text{CFL} \leq 1$ also for SPELT due to some data structure limitations in HOMME.

Remark 4.1. Perhaps the major advantages of semi-Lagrangian schemes are to overcome strong CFL restrictions. However, for a parallelization we think that there is the need to find a good balance between higher CFL and the bandwidth for communication. We note that especially for CSLAM a $\text{CFL} > 1$ is rather complicated to realize in HOMME due to the 5×5 reconstruction stencil.

In Figure 5 we see the data situation for the spectral element scheme, SPELT, and CSLAM. The spectral element is marked as a bold square. The different size in Figures 5(a)–(c) is because of visualization reasons. Figure 5(a) shows the situation for the spectral element with 5×5 Gaussian-Legendre-Lobatto (GLL) points, which is comparable with 4×4 cells for the semi-Lagrangian schemes. Note that there is not a classical halo zone necessary, instead the filled circles are communicated such that the continuous constraint for the spectral element scheme is fulfilled. In Figure 5(b) the filled squares and circles represent the points for the SPELT reconstruction with 4×4 cells for a spectral element. Note that the reconstruction is globally continuous, so only the black squares and circles are in the halo zone (the points on the boundary are already on the element). The depth of

the halo zone is one cell which allows a $CFL \leq 1$. For CSLAM the situation is illustrated in Figure 5(c). For a $CFL \leq 1$ we need a halo zone with four cells. That means for a spectral element we store 25 data in the spectral element and 16 points are sent from a different processor (here some kind of averaging with the local points on the boundary). SPELT stores 162 and 88 points are sent from a different processor. For CSLAM we store 144 data and 128 points come from a different processor.

Let us consider a resolution of $0.25^{\circ a}$ on the equator, 100 tracers and 26 levels (the vertical discretization is divided in levels, i.e., we basically consider here 26 horizontal discretizations) in the climate model CAM. The following (overall) memory would be allocated for the standalone tracer transport schemes; 24 101 MB (SE), 162 924 MB (SPELT) and 138 823 MB (CSLAM). The memory usage and the halo zone values for both semi-Lagrangian schemes are similar. However, compared to the spectral element scheme it is huge. Nevertheless, if you divide this number by the number of processors, the usage per processor is acceptable, since for this high resolution you would likely need more than 1000 processors.

Remark 4.2. To apply the positivity preservation approach of [17] or any other limiter for SPELT we need information from a cell patch. If we allow a deeper halo zone (i.e., more values for communication) we could implement a positivity limiter without an additional communication. However, we think this contradicts the idea to develop a scheme for non-uniform or unstructured grids and allow therefore a second communication. This is not true for the QMSL filter described above, which can easily be implemented locally without any extra communication. For CSLAM, however, the novelty of having only one communication per time step is valid for both, using the scheme with and without monotonicity preserving filtering (both with the same amount of communication data).

Remark 4.3. The HOMME algorithm of CSLAM in [13] and thus the data structure can be used also for the flux form CSLAM described in [25]. The only difference is that one has to calculate four areas for the fluxes, see also Remark 3.1 and Figure 3. Instead of finding the intersection of the departure grid to the Eulerian grid, one has to find intersection to the Eulerian grid of the flux areas. In particular, that means one has to do the search four times instead of only one time. The reconstruction and the filter are exactly the same.

^aThe highest resolution for CAM-SE is up to 0.125° , which is considered as a very high resolution for climate simulation.

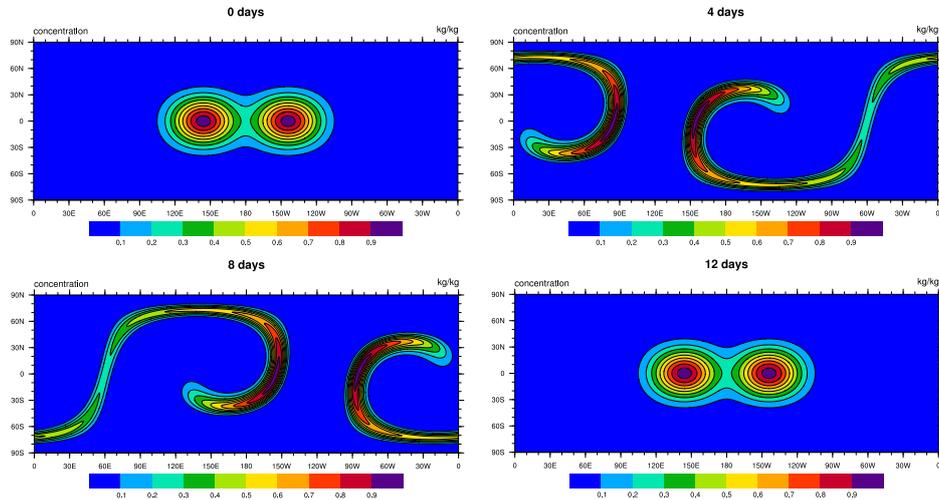


Figure 6. Deformational test calculated with SPELT in HOMME for a mesh size (cell size) of 0.28° at the equator and with no monotonicity constraint. The time step is chosen such that $\text{CFL} \leq 0.8$. The two Gaussian distributions return to their initial position after 12 days. The figure is taken from [12].

4.3. Error and strong scalability analysis for SPELT and CSLAM

In the following we consider the smooth deformational-flow test introduced in [26]. For the initial distribution we use two Gaussian hills which will be highly deformed through a non-divergent wind field. After 12 days the flow deforms back to the smooth exact solution. Therefore, the numerical solution after 12 days is suitable for error and convergence studies, e.g., in the L^1 , L^2 , and L^∞ error norms. For all test cases in this section we choose the time step in such a way that the maximal CFL is always ≤ 0.8 (fixed) for all resolutions. For a general description of this test case we refer to [26]. There, we can find an accurate trajectory computation algorithm (to find the departure points), which is a combination of a Taylor series expansion and splitting the trajectory into segments. We can apply this algorithm in this section since the velocities are analytically known and therefore available in any point on the sphere. In Section 5 we will discuss a departure grid algorithm for practical applications, i.e., where the velocities are known only on certain points.

In Figure 6 we show the time evolution of our chosen test example for a mesh size (cell size) of 0.28° at the equator. This corresponds to 320×320 cells on one cube face. The convergence plots in Figure 7 show the normalized L^1 , L^2 , and L^∞ error norms and confirm for both schemes, that

Scalable, conservative, multi-tracer efficient semi-Lagrangian schemes

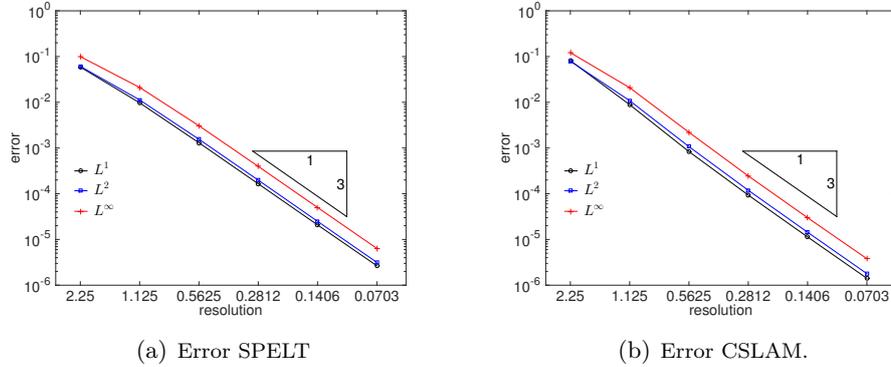


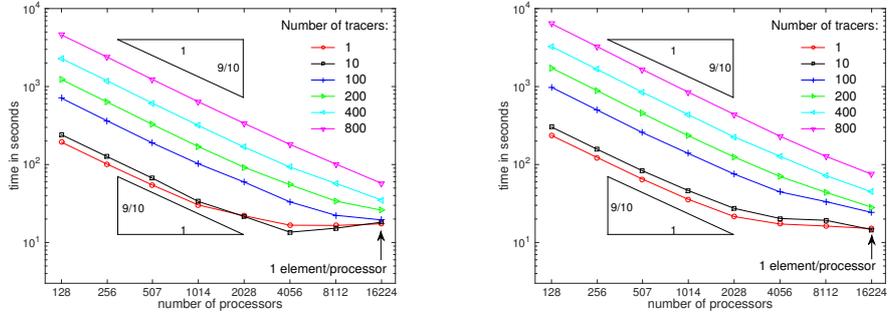
Figure 7. The plots show the convergence order for different normalized error norms L^1 , L^2 , and L^∞ with respect to the degree resolution at the equator. As a test example we use the deformational flow test, compare Figure 6.

they are third order accurate for a sufficient smooth initial field. The mesh size (cell size) for the last run is 0.0703° at the equator, which corresponds to 1280×1280 cells on a cube face.

A more detailed error behavior discussion for both finite-volume schemes, SPELT and CSLAM in HOMME, can be found in [12]. In terms of accuracy both schemes lead approximately to the same relative error. However, in [12] we got the remarkable findings that for smooth fields CSLAM seems to be slightly more accurate than SPELT, but for non-smooth or quasi-smooth fields SPELT is slightly more accurate than CSLAM.

Note that both semi-Lagrangian schemes are built up on the spectral element grid of HOMME. Thus we expect to benefit from the excellent scalability of the horizontal decomposition of the HOMME model [2]. Figure 8 shows some (strong) scalability results on Yellowstone, which is the new (2013) petascale computing resource in the NCAR-Wyoming Supercomputing Center. Yellowstone is a 1.5-petaflops high-performance computing system with 2.6-GHz Intel Xeon Sandy Bridge processors and 2 GB memory/processor^b. A standard user can allocate up to 16384 processors. For our test we run our benchmark test using a problem size of 16224 spectral elements (equidistant mesh on the cubed-sphere), which allows us to test our algorithm for one spectral element per processor. Each of these elements are divided in 4×4 cells (equidistant finite volume grid) as in Figure 4(b). This results in 2704 spectral elements per cube face or 43264 cells. To ensure $\text{CFL} \leq 1$ we allow a time step of 600 seconds. For all tests we disable I/O. Figure 8 shows the strong scalability with respect to the num-

^b<https://www2.cisl.ucar.edu/resources/yellowstone>



(a) Spelt with QMSL filter and positivity option.

(b) CSLAM with monotonic option.

Figure 8. Both figures, taken from [12], show strong scalability for different numbers of tracers on Yellowstone (2.6-GHz Intel Xeon Sandy Bridge processors). The test example is the deformational-flow test similar as in Figure 6. We choose the input data in such a way that the right most marker on each line with 16224 processors runs on the minimum of one spectral element per processor or 16 unknowns per tracer and processor. Panel (a) shows the results for SPELT combined with a QMSL filter and an additional positivity limiter. This configuration results in two communication per time step in our implementation, see also Remark 4.2. Panel (b) shows CSLAM scalability with the monotonicity option turned on.

ber of processor. The right most marker on each line is for 16224 processors and corresponds to one spectral element per processor or 16 unknowns per processor and per tracer.

In Figure 8(a) we see the strong scalability results for SPELT with a QMSL filter and the positivity preserving option turned on. This results in a second communication, see Remark 4.2. The algorithm does not scale below 4 elements per processor and 1 – 10 tracers but scales again for more than 100 tracers. We compare our results for SPELT with the strong scalability of CSLAM on Yellowstone shown in Figure 8(b). We stress that the scalability of CSLAM with one communication per time step is similar to the SPELT algorithm with two communications in Figure 8(a) for a low number of tracers and element per processor, which is due to the higher communication load of the scheme (CSLAM gets 128 and SPELT 88 values from different processors). Note that the SPELT algorithm with the positivity preserving option (and therefore two communications per time step) is faster for more tracers or more elements per processor. Increasing the number of tracers (more than 100) leads to perfect strong scaling (0.9) also for 1 element per processor for both schemes. All runs show also the multi-tracer efficiency for SPELT and CSLAM. More details can be found in [12]. We also refer to the results in [13] for the performance of CSLAM on an outdated IBM Blue

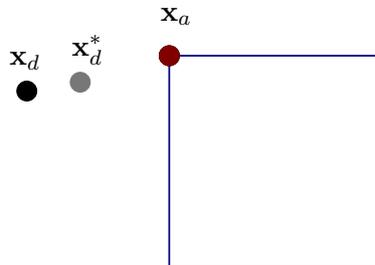


Figure 9. One spectral element with one arrival point \mathbf{x}_a . With \mathbf{x}_d^* we denote the predicted departure point from the first step of our departure point calculation. In a second step we get the final (approximated) departure point \mathbf{x}_d . Note that from a geometrical viewpoint \mathbf{x}_d may come from another element - which could be on another processor.

Gene/L system with a much slower CPU per processor. This machine had 4096 processors and even the configuration of one tracer and one element per processor led to a good strong scalability.

5. A departure grid algorithm for HOMME

In Section 4 we used a highly accurate algorithm to find the departure point. There, a benchmark test was applied where the velocity field is known (analytically) in any point and at any time. This is not met in practice. Therefore, we introduce in this section a departure algorithm with higher accuracy than a simple linear determination and apply our approach in Section 6 on a baroclinic instability test. It consists of a classical Runge-Kutta and Taylor expansion ansatz. The method is based on the one presented in [27], where a Taylor series expansion and an approximation argument of the total derivative leads also to an economical determination of departure points with higher accuracy. Using derivative information from the spectral element method, we can compute a second-order accurate estimation of the departure point with no additional communication or time-stepping along backward trajectories. That means, we need only the velocities at one point but on different time levels. We also refer to the work [28], where the Taylor series method [27] is compared with a Runge Kutta scheme. There, the velocities at intermediate time levels for forward trajectories are estimated by linear or cubic interpolation.

Remark 5.1. In this work we focus on semi-Lagrangian schemes for passive tracer transport. Thus, the dynamics, and in particular the velocities, are solved by another scheme (here the spectral element method). That means the velocities $\mathbf{u} = (u, v)$, see e.g. (1), are given by a high order approximation on each element at time t and $t + \Delta t$.

In the following \mathbf{x}_d denotes the departure point which moves to the arrival point \mathbf{x}_a during the time $[t, t + \Delta t]$ with the time step Δt ; see Figure 9. For the coupling to the spectral element grid, where the velocities are given at \mathbf{x}_a , we need an economical determination of \mathbf{x}_d . With a horizontal domain decomposition, the spectral element containing the point \mathbf{x}_a and the spectral element containing the point \mathbf{x}_d may not be on the same processor. In order to minimize communication between these processors, we look for an estimate of \mathbf{x}_d that can be computed using only information given at \mathbf{x}_a .

The classical predictor corrector scheme, which belongs to the class of Runge-Kutta schemes, reads for backward trajectory: first we get a predictor for the departure point \mathbf{x}_d ;

$$\mathbf{x}_d^* = \mathbf{x}_a - \Delta t \mathbf{u}(\mathbf{x}_a, t + \Delta t).$$

Then the corrector and finally the departure point is given by

$$\mathbf{x}_d = \mathbf{x}_a - \Delta t \frac{1}{2} (\mathbf{u}(\mathbf{x}_d^*, t) + \mathbf{u}(\mathbf{x}_a, t + \Delta t)).$$

Note that the velocity $\mathbf{u}(\mathbf{x}_d^*, t)$ may be calculated using velocity data from the spectral element containing the point \mathbf{x}_d^* , but as described above we instead estimate this value using only data at the point \mathbf{x}_a . By Taylor series we write

$$\begin{aligned} \mathbf{u}(\mathbf{x}_d^*, t) &= \mathbf{u}(\mathbf{x}_a, t) + (\mathbf{x}_d^* - \mathbf{x}_a) \nabla \mathbf{u}(\mathbf{x}_a, t) \\ &= \mathbf{u}(\mathbf{x}_a, t) - \Delta t \mathbf{u}(\mathbf{x}_a, t + \Delta t) \nabla \mathbf{u}(\mathbf{x}_a, t). \end{aligned}$$

Finally, we can calculate the departure points in terms of the arrival velocity $\mathbf{u}(\mathbf{x}_a, \cdot)$

$$(8) \quad \mathbf{x}_d = \mathbf{x}_a - \Delta t \left(\frac{\mathbf{u}(\mathbf{x}_a, t) + \mathbf{u}(\mathbf{x}_a, t + \Delta t)}{2} - \frac{\Delta t}{2} \mathbf{u}(\mathbf{x}_a, t + \Delta t) \nabla \mathbf{u}(\mathbf{x}_a, t) \right)$$

Note that approximation (8) is second order accurate in space and time.

6. Comparison: Spectral Elements (SE) Advection Scheme versus SPELT and CSLAM

In this section we show the first results using SPELT or CSLAM for the passive tracer transport. HOMME and therefore CAM-SE uses a conventional vector-invariant form of the moist primitive equations. For the vertical discretization it uses the hybrid η pressure vertical coordinate system. For more details we refer to [2]. The model is split into a horizontal (two dimensional) and a vertical discretization. The vertical discretization

Scalable, conservative, multi-tracer efficient semi-Lagrangian schemes

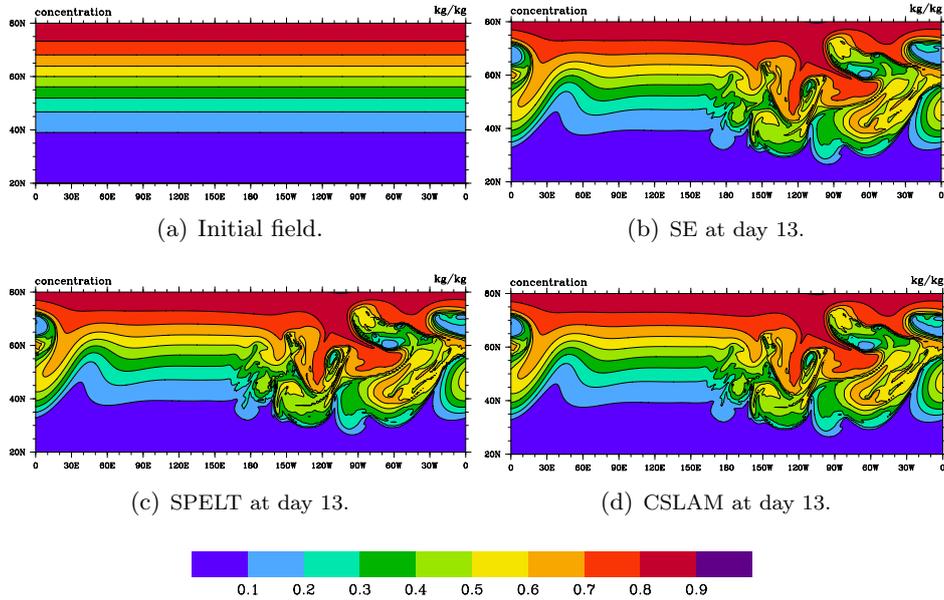


Figure 10. Figure (a) shows the initial field, the passive tracer, for the baroclinic instability test in the northern hemisphere, where the instability develops. The test uses the configuration described in [30]. Figures (b), (c), and (d) show the tracer field after 13 days of advection. The excluded domain does not show any deforming of the initial field. Thus we highlight only the region of interest.

works in the current configuration on 26 reference levels with a vertically Lagrangian approach [29]. That means the advection schemes used in horizontal directions on the floating Lagrangian levels are occasionally remapped back (vertical one dimensional) to the reference levels.

In this section we focus on the horizontal discretization. For that the mass continuity equation reads for the dynamics

$$(9) \quad \frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \mathbf{u}) = 0$$

with the air density ρ and the velocity $\mathbf{u} = (u, v)$. In the following, (9) is always approximated by the spectral element scheme and is part of the dynamics. In particular, the velocity $\mathbf{u} = (u, v)$ is known for passive transport. The passive tracers are described by the conservation equation

$$(10) \quad \frac{\partial(\rho\phi)}{\partial t} + \nabla \cdot (\rho\phi\mathbf{u}) = 0,$$

where ϕ is the tracer concentration. In the following we use an idealized baroclinic instability flow benchmark test, which is solved by the dynamics, i.e., the spectral element scheme. A detailed description can be found

in [30]. The spectral element transport, SPELT, and CSLAM are used to find a discrete solution for (10), i.e, we use only these schemes for passive transport. To calculate the departure grid for the semi-Lagrangian schemes we apply the algorithm described in Section 5. We use a high resolution of 120×120 spectral elements per cube face with 5×5 Gaussian points for the spectral element and 4×4 cells per spectral element for SPELT and CSLAM. This means a very high resolution of 0.1875° at the equator. We choose the time in such a way that the spectral element scheme has a CFL= 0.27 and SPELT and CSLAM CFL= 1. Thus all schemes with their monotone options are stable and run with the maximal CFL, which is available in HOMME.

We start with the tracer initial field shown in Figure 10(a). The instability starts to grow in the northern hemisphere. In Figures 10(b),(c), and (d) we see the field after 13 days for the passive transport with the spectral element method, SPELT, and CSLAM. We stress that until day 13 the instability happens only between the latitudes 20N and 80N. In other words, the excluded domain does not show any deforming of the initial field, and we want to highlight only the region of interest. Note that there is no analytical solution available for this test. Thus a precise error discussion might be a difficult task. Note that especially in the instability region, all three schemes differ a little bit, i.e., the difference field is small. However, we do not know which scheme convergences best and presents the right wave behavior. Thus it would be speculating to use one as a reference solution. Nevertheless, more tests have to be done on this type of coupling which is beyond the scope of the paper.

Remark 6.1. It is an ongoing research to provide a consistent numerical coupling between the dynamics (9) and the passive tracer mass transport (10) in the sense of [31]. Consistency means, that if $\phi = 1$ in (10) then the tracer scheme should match the results produced by the scheme used for (9). Note that the quantities that should be conserved are ρ and $\rho\phi$, respectively. A possible consistent coupling with the remapping scheme CSLAM is introduced in [26, Appendix B] if both, (9) and (10), are approximated by CSLAM. However, in our model we approximate (9) by a spectral element scheme, which is in particular not a remapping scheme. Thus we did not succeed to find an efficient solution for the consistent coupling between these two different schemes. For numerical schemes in flux form consistent couplings are proposed in [31–33] to mention only a few but not all. Note that the spectral element transport scheme is based on a Galerkin ansatz and thus contradicts the philosophy of an explicit flux approximation. However, it has been shown that it is possible to calculate

Scalable, conservative, multi-tracer efficient semi-Lagrangian schemes

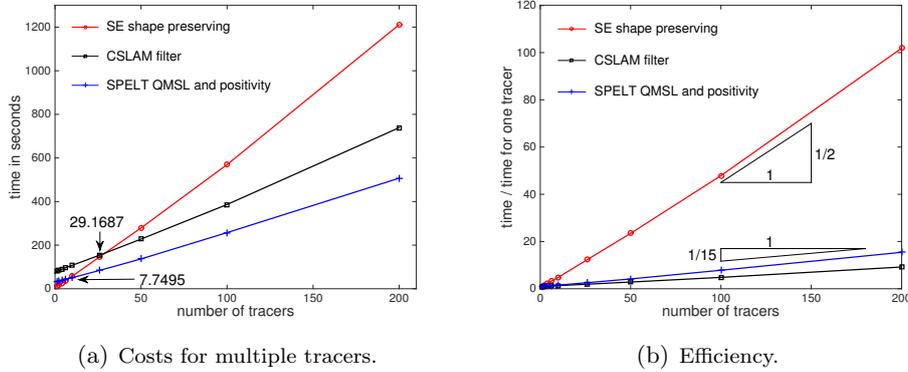


Figure 11. Comparison of computational costs of the spectral element transport scheme, SPELT, and CSLAM for the horizontal discretizations of the passive tracer transport equation in the baroclinic instability test on 1024 processors (left). Right: the individual efficiency of the schemes. The slope indicates the additional costs for adding one tracer compared to one tracer of the corresponding scheme. E.g., for the spectral element transport we need 1/2 of the time of the first tracer, to add a second tracer. Note that despite the illustration here, SPELT is more efficient than CSLAM. The figure shows in fact, how the single scheme performs with respect to the first tracer.

fluxes from the spectral element solution. Thus there is the hope to formulate a consistent coupling between the spectral element dynamics and flux form semi-Lagrangian schemes such as SPELT [12] and the flux form version of CSLAM [25] in the sense of [31–33]. However, this is beyond the scope of this paper.

In Figure 11 we show the performance of the schemes. As a test we use the baroclinic instability test described above with a resolution of 0.75° at the equator, 15 days and a spectral element $CFL \leq 0.27$. For SPELT and CSLAM we allow $CFL \leq 1$. Additionally, we run the schemes in shape preserving mode. Thus all three passive tracer transport schemes run stable with the maximal CFL, which is currently available in HOMME. For this test we allocate 1024 processors. We consider the overall time for the horizontal discretizations (26 levels), including the calculation of the departure grids. In Figure 11(a) we see the computational time of all the three transport schemes. We observe that the spectral element scheme is very efficient for less than 10 tracers but becomes very computational expensive the more tracer we add. For both semi-Lagrangian schemes it is the other way around. The first tracer is rather expensive because of the pre calculation of certain quantities. However, both, SPELT and CSLAM, benefit from their multi-tracer efficiency the more tracers we add, i.e., due to their

construction, they can reuse quantities for additional tracers as described in Section 2 and Section 3, respectively. In particular, SPELT is cheaper for more than 7.7495 tracers than the spectral element scheme and CSLAM for more than 29.1687 tracers. Note that SPELT is less expensive for one tracer compared to CSLAM and performs better the more tracer we add. In Figure 11(b) we plot an efficiency plot. We see the performance for each individual scheme with respect to the time for one tracer of the corresponding scheme. For example, if we add one tracer, the spectral element scheme needs half of the computational time of the first tracer, aso. In this figure, it seems that CSLAM performs slightly better than SPELT. However, this plot shows only the performance of the individual scheme, i.e., we compare the scheme with itself.

7. Conclusions

In this work we presented two semi-Lagrangian schemes, SPELT and CSLAM, which are mass conservative and multi-tracer efficient. In particular, we discussed several issues on multiple processor platforms and showed the convergence and strong scalability in HOMME on a standard benchmark test and state-of-the-art supercomputer. We remark that SPELT is capable of handling arbitrary unstructured quads and has the potential to be used for variable resolution cubed-sphere grids for regional climate modeling in CAM-SE. This is not obvious for CSLAM. To find the departure points, which are needed for both schemes, we introduced an economical algorithm designed for the spectral element grid. Finally, the comparison to the default spectral element tracer transport scheme in HOMME shows that both semi-Lagrangian schemes are highly multi-tracer efficient. Two major ingredients lead to that; semi-Lagrangian schemes allow longer time steps than explicit schemes on an Eulerian mesh, i.e., currently in HOMME the difference is $CFL \leq 1$ to $CFL \leq 0.27$. Second, SPELT and CSLAM allow to reuse information, which makes them multi-tracer efficient and reduces the number of communication per time step. On the other hand, the price we pay for longer time steps for tracer transport and a low number of communication is the increase of memory per processor. It is an ongoing work to provide a consistent coupling between the spectral element dynamics and the semi-Lagrangian schemes for passive tracer transport. However, first results on the coupling have been shown in this work. There is no intention to work on a 3D implementation in the near future since the current dynamical core HOMME for climate modeling uses a dimension splitting approach and simply does not provide the data structure for a fully 3D model. In principle, both schemes could be extended to 3D. However, it

seems to be almost impossible to do it for CSLAM since the search for overlap regions might be hard to realize in 3D. SPELT, however, relies on pointwise search and is therefore also suitable for 3D. Due to the simple formulation of the SPELT scheme it might be also appropriate for numerical weather prediction models.

REFERENCES

1. D. L. Williamson, The evolution of dynamical cores for global atmospheric models, *Journal of the Meteorological Society of Japan*, vol. 85, pp. 241–269, 2007.
2. J. Dennis, J. Edwards, K. Evans, O. Guba, P. Lauritzen, A. Mirin, A. St-Cyr, M. Taylor, and P. Worley, CAM-SE: A scalable spectral element dynamical core for the Community Atmosphere Model, *International Journal of High Performance Computing Applications*, vol. 26, no. 1, pp. 74–89, 2012.
3. M. Taylor, J. Edwards, S. Thomas, and R. Nair, A mass and energy conserving spectral element atmospheric dynamical core on the cubed-sphere, *Journal of Physics: Conference Series*, vol. 78, 2007.
4. F. X. Giraldo, Lagrange-Galerkin methods on spherical geodesic grids: the shallow water equations, *Journal of Computational Physics*, vol. 160, pp. 336–368, 2000.
5. R. D. Nair, S. J. Thomas, and R. D. Loft, A discontinuous Galerkin global shallow water model, *Monthly Weather Review*, vol. 133, pp. 876–888, 2005.
6. J. M. Dennis, A. Fournier, W. F. Spitz, A. St-Cyr, M. A. Taylor, S. J. Thomas, and H. M. Tufo, High resolution mesh convergence properties and parallel efficiency of a spectral element atmospheric dynamical core, *International Journal of High Performance Computing Applications*, vol. 19, pp. 225–235, 2005.
7. R. Sadourny, Conservative finite-difference approximations of the primitive equations on quasi-uniform spherical grids, *Monthly Weather Review*, vol. 100, pp. 136–144, 1972.
8. O. Guba, M. A. Taylor, and A. St-Cyr, Optimization-based limiters for the spectral element method, *Journal of Computational Physics*, vol. 267, pp. 176–195, 2014.
9. M. Zerroukat and T. Allen, A three-dimensional monotone and conservative semi-Lagrangian scheme (SLICE-3D) for transport problems, *Quarterly Journal of the Royal Meteorological Society*, vol. 138, pp. 1640–1651, 2012.
10. R. D. Nair, J. S. Scroggs, and F. H. M. Semazzi, Efficient conservative global transport schemes for climate and atmospheric chemistry models,

- Monthly Weather Review*, vol. 130, no. 8, pp. 2059–2073, 2002.
11. P. Bochev, D. Ridzal, and K. Peterson, Optimization-based remap and transport: A divide and conquer strategy for feature-preserving discretizations, *Journal of Computational Physics*, vol. 257, pp. 1113–1139, 2014.
 12. C. Erath and R. D. Nair, A conservative multi-tracer transport scheme for spectral-element spherical grids, *Journal of Computational Physics*, vol. 256, no. C, pp. 118–134, 2014.
 13. C. Erath, P. H. Lauritzen, J. H. Garcia, and H. M. Tufo, Integrating a scalable and efficient semi-Lagrangian multi-tracer transport scheme in HOMME, *Procedia Computer Science*, vol. 9, pp. 994–1003, 2012.
 14. C. G. Chen, F. Xiao, X. L. Li, and Y. Yang, A multi-moment transport model on cubed-sphere grid, *International Journal for Numerical Methods in Fluids*, vol. 67, no. 12, pp. 1993–2014, 2011.
 15. S. T. Zalesak, Fully multidimensional flux-corrected transport algorithms for fluids, *Journal of Computational Physics*, vol. 31, pp. 335–362, 1979.
 16. C. Chen and F. Xiao, Shallow water model on cubed-sphere by multi-moment finite volume method, *Journal of Computational Physics*, vol. 227, pp. 5019–5044, 2008.
 17. D. R. Durran, *Numerical Methods for Fluid Dynamics with Applications to Geophysics*. Springer, 2010.
 18. P. H. Lauritzen, R. D. Nair, and P. A. Ullrich, A conservative semi-Lagrangian multi-tracer transport scheme (CSLAM) on the cubed-sphere grid, *Journal of Computational Physics*, vol. 229, no. 5, pp. 1401–1424, 2010.
 19. J. K. Dukowicz and J. R. Baumgardner, Incremental Remapping as a Transport/Advection Algorithm, *Journal of Computational Physics*, vol. 160, pp. 318–335, 2000.
 20. J. K. Dukowicz, Conservative rezoning (remapping) for general quadrilateral meshes, *Journal of Computational Physics*, vol. 54, pp. 411–424, 1984.
 21. J. K. Dukowicz and J. Kodis, Accurate conservative remapping (rezoning) for arbitrary lagrangian-eulerian computations, *SIAM Journal on Scientific and Statistical Computing*, vol. 8, no. 3, pp. 305–321, 1987.
 22. P. A. Ullrich, P. H. Lauritzen, and C. Jablonowski, Geometrically Exact Conservative Remapping (GECORE): Regular latitude-longitude and cubed-sphere grids, *Monthly Weather Review*, vol. 137, pp. 1721–1741, 2009.
 23. C. Erath, P. Lauritzen, and H. Tufo, On Mass Conservation in High-Order High-Resolution Rigorous Remapping Schemes on the Sphere,

Scalable, conservative, multi-tracer efficient semi-Lagrangian schemes

- Monthly Weather Review*, vol. 141, no. 6, pp. 2128–2133, 2013.
24. T. J. Barth and D. C. Jespersen, The design and application of upwind schemes on unstructured meshes, *27th Aerospace sciences meeting*, vol. 89, no. 89-0366, 1989.
 25. L. M. Harris, P. H. Lauritzen, and R. Mittal, A flux-form version of the conservative semi-Lagrangian multi-tracer transport scheme (CSLAM) on the cubed sphere grid, *Journal of Computational Physics*, vol. 230, no. 4, pp. 1215–1237, 2011.
 26. R. D. Nair and P. H. Lauritzen, A class of deformational flow test cases for linear transport problems on the sphere, *Journal of Computational Physics*, vol. 229, p. 8868, 2010.
 27. J. L. McGregor, Economical Determination of Departure Points for Semi-Lagrangian Models, *Monthly Weather Review*, vol. 121, no. 6, pp. 221–230, 1993.
 28. R. D. Nair, J. S. Scroggs, and F. H. M. Semazzi, A forward-trajectory global semi-lagrangian transport scheme, *Journal of Computational Physics*, vol. 190, pp. 275–294, 2003.
 29. S.-J. Lin, A vertically lagrangian finite-volume dynamical core for global models, *Monthly Weather Review*, vol. 132, pp. 2293–2397, 2004.
 30. C. Jablonowski and D. L. Williamson, A baroclinic instability test case for atmospheric model dynamical cores, *Quarterly Journal of the Royal Meteorological Society*, vol. 132, pp. 2943–2975, 2006.
 31. S.-J. Lin and R. B. Rood, Multidimensional flux-form semi-Lagrangian transport schemes, *Monthly Weather Review*, vol. 124, no. 9, pp. 2046–2070, 1996.
 32. C. Schär and P. K. Smolarkiewicz, A synchronous and iterative flux-correction formalism for coupled transport equations, *Journal of Computational Physics*, vol. 128, pp. 101–120, 1996.
 33. E. S. Gross, L. Bonaventura, and G. Rosatti, Consistency with continuity in conservative advection schemes for free-surface models, *International Journal for Numerical Methods in Fluids*, vol. 38, pp. 307–327, 2002.