Reviewers 1 and 2 and Associate Editor:

Thank you for your reviews. In previous responses, I addressed a number of questions and technical comments:

- to reviewer 1: https://doi.org/10.5194/gmd-2021-296-AC2;
- to reviewer 2: https://doi.org/10.5194/gmd-2021-296-AC3.

In this response accompanying the submission of a revised manuscript, I respond to points raised about the text and figures.

I carried out a major revision to the original manuscript. First, as suggested, I have divided the original manuscript into two. Most, but not all, of the material from Sects. 4 and 5 are in this revised manuscript. The material from Sects. 2 and 3 is used in another manuscript.

Second, to clarify material, I have added much more background information. In the new manuscript, Sect. 1 discusses very high-level matters. Section 2 provides details about all the algorithmic subcomponents that are then assembled into the overall method in Sect. 3. Each of the subsections of Sect. 2 has two parts: high-level explanations followed by mathematical details. In a first reading, the mathematical details can be skipped.

Third, Sect. 4 contains fewer figures and more comparisons to other methods in the text.

Fourth, Appendix A summarizes the Islet bases for completeness, while omitting almost all the details of their derivation; details are now in a second manuscript.

Fifth, in addition to removing material about the derivation of the Islet bases, I have also removed computer implementation details and performance figures; these are in the second manuscript. However, to provide clear information about the high performance of the Islet method, I have added new text regarding performance in the recently released E3SM version 2 early in the introduction (lines 18–26).

In the following, the reviewer's comments are italicized. Revised text from the manuscript is blue.

RC1: *The paper presents a new set of basis functions for semi-Lagrangian advection method in spectral element models. The presented method is novel and its performance in atmospheric test cases is encouraging. The paper therefore warrants to be published. The main drawback of the paper is that it is quite tedious to read. This paper essentially presents two things: (1) a novel optmized basis functions for 1D semi-Lagrangian advection schemes, and (2) its implementation in an atmospheric model using three different (sub-element) grids. Both of these are quite complex topics, and their discussion is intertwined (e.g. in section 2) and the reader is easily lost in details. The overall presentation should be improved before the paper can be accepted for publication. Considering the amount of work, I recommend a major review.*

Thanks. In reworking the presentation, I have tried to reduce the tedium by making each subsection in Sect. 2 have a high-level discussion followed by mathematical details that can be skipped. Section 3 now has no low-level details; instead, references to the mathematical details in subsections of Sect. 2 are provided.

RC1: *For clarity, I suggest that you clearly define the two considered problems, 1D SL advection method, and the implementation in 3-grid atmospheric models, already in the introduction. The introduction is now quite short and actually does not mention many important concepts relevant for the paper.*

Thanks. Sections 1 and 2 are now an extended introduction. In particular, both the 3-grid model and the advection method are discussed at length, first at a high level and referenced to

Figs. 1 and 2, then with mathematical details. Figure 2 is new and illustrates the interpolation SL method on spectral elements.

RC1: *Section 2 is rather difficult to follow, consider revising. I suggest to start by defining the 1D discretization with N elements, the interpolant functions within each element, and the properties of the interpolant functions (e.g. basis functions, continuity and symmetry). Presently, the interpolant functions first appear only in section 3.1. A figure could clarify the concepts, including the source and target elements/nodes. I would also define a symbol for the basis functions themselves, instead of using L from section 3.1 (L is the interpolant function itself). The discussion of stability becomes comprehensible only after the discretization has been introduced.*

Thanks. The stability material will appear in a second paper. I have attempted to build up the rest of the ideas piece by piece in Sect. 2. The new Fig. 2 illustrates target and source elements, GLL nodes on the dynamics and tracer grids, and 1D and 2D basis functions.

RC1: *As I understand, the 3 axioms of advection methods are: (1) global conservation, (2) preservation of constant tracers (sometimes called local conservation, or tracer consistency), and (3) monotonicity (i.e. no spurious overshoots appear). These properties apply to both the SL tracer advection scheme and the remap operators between grids. In section 1.2 these concepts seem to be mixed and referred to by different names (plausibly due to historical reasons; the so called "property preservation" is just a combination of 1, 2, and 3). Consider revising.*

Local mass conservation is not the same as mass-tracer consistency. A solution can be locally mass conserving but not mass-tracer consistent, as well as the opposite. Mass-tracer consistency means that the transport solver and the dynamics solver produce or use the same air density field ρ to machine precision.

I have rewritten that section as follows:

(48–57) An approximate numerical solution for q of Eq. 4 is said to be *property preserving* if (possibly just a subset of) properties that hold for the exact solution also hold for the approximate one. Equation 5 implies that advection cannot introduce new extrema in the mixing ratio; advection is said to be *shape preserving*. Equation 2 with $f = 0$ implies the global mass is conserved. Although the focus of this article is not the continuity equation, we note that the Lagrangian form of the continuity equation, Eq. B4 in Appendix B, implies that the total mass in a Lagrangian parcel, which is a parcel of fluid that moves with the flow, is constant. A final property that is a special case of the shape preserving property relates to coupling a solver for Eq. 4 to a dynamics solver: mass-tracer consistency. This property means that if q is constant in space at time t_0, then it remains constant in space at every other time. In other words, the dynamics solver and transport solver use the same air density. The methods in this article conserve global mass, do not introduce new nodal extrema, and provide mass-tracer consistency when coupled to a dynamics solver.

Regarding local mass conservation, I now write:

(167–170) Local mass conservation means that one can identify numerical, possibly Lagrangian, fluid parcels on the grid that have constant tracer mass. *Local* is in contrast to *global* mass conservation; the latter means that the mass of the tracer fluid is conserved over the whole domain but not necessarily in any identifiable parcels smaller than the domain.

Finally, In the new Sect. 2.2, which I do not quote here because of its length, I write the specific property preservation problem the Islet method solves in detail, thus resolving any ambiguities that may arise from different terminology.

RC1: *Figures 20 and 21 shows larger diffusion for shorter time steps. Could you elaborate on this? Is the traditional SL property that longer time steps reduce diffusion (while the solution can degrade in other metrics)? How would one choose the right time step in practice? In the case of Eulerian transport it is easy: take the maximum stable one, and you are guaranteed to satisfy all the necessary properties.*

To address the first part of this question, I added the following.

(549–552) Essentially all SL methods, particularly when given exact trajectory data, exhibit greater error with smaller time step, e.g., CSLAM in TR14. This is because the only source of error, given exact trajectories, is the remap error. Smaller time steps correspond to more remaps to reach a fixed simulation time.

I answered the second part of this question in the original response to reviewer 1. I have not added material to the paper to answer this part because the answer is complex and somewhat orthogonal to the primary material. All SL transport methods must contend with the matter of time step. In the standard-resolution EAM version 2 configurations, the time step is limited by the physics time step to six times the dynamics time step.

RC1: *Throughout the manuscript the authors use the terms "mixing ratio" and "tracer" for the advected quantity q_i, seemingly interchangeably. For the sake of clarity I would prefer just to use "tracer".*

Thanks. A tracer is a trace species, such as CO_2. A mixing ratio is a dimensionless quantity. In meteorology the mixing ratio is typically the mass mixing ratio, e.g., grams of a species per kilograms of wet air. Thus, a tracer is not the same as a mixing ratio; the tracer is the substance whose quantity can be given by a mixing ratio or a density.

To address this point, I have added the following:

(33–34) The tracer transport equation in continuity form and with a source term for a tracer mixing ratio q and corresponding tracer density ρq is...

I have also been careful to use *tracer mixing ratio* (q), *tracer density* (ρq) and *air density* (ρ) consistently throughout the manuscript.

RC1: *l35: what is the definition of "local conservation" here?*

I have added the following text:

(167–170) Local mass conservation means that one can identify numerical, possibly Lagrangian, fluid parcels on the grid that have constant tracer mass. *Local* is in contrast to *global* mass conservation; the latter means that the mass of the tracer fluid is conserved over the whole domain but not necessarily in any identifiable parcels smaller than the domain.

RC1: *l155: "This structure arises as follows. Consider a continuous discretization using a nodal n_p-basis, $n_p = d+1$, with n_p the number of nodes. The grid has N elements. Each row of the space-time matrix corresponds to a target node." This description is too brief to be understandable, please elaborate. This is my interpretation of the discretization: The 1D domain is divided into N elements. The solution in each element is approximated by a continuous function, defined by n_p basis functions. Thus a function f in an element e can be written as $f_e(x) = \sum_{i=1}^{n_p} f_i \psi_i(x)$ where ψ_i and f_i denote the i-th basis function and its corresponding coefficient. Each basis function is associated with a node x_i within the element; The basis is Lagrangian (a.k.a. nodal), i.e. $f_i(x_i) = 1$. Furthermore, the discretization of the function is continous across element interfaces, implying that the neighboring elements share a (exactly one?) basis function. Furthermore, the basis is assumed to be symmetric about the center point of the element.*

This passage has been removed, as the topic will appear in a second paper. However, I have added text following your suggestions in Sect. 2 to clarify relationships among elements, nodes, and bases.

RC1: *l156: "Each row of the space-time matrix corresponds to a target node." You should define the space-time matrix for this statement to be comprehensible.*

I have removed references to the space-time matrix in this manuscript, as the matrix is needed only for stability analysis, which will appear in a companion paper.

RC1: *l177: Have you defined the basis to be symmetric somewhere?*

In the revised manuscript, I write that the basis is symmetric in Sect. 2.1.2 (background material) and Appendix A (precise description of the Islet bases):

(186–187) Each basis is a *nodal* basis: a basis function has value 1 at one node and 0 at every other node. Thus, each basis function is associated with a node. For example, in Fig. 2(b), the blue basis function is associated with the third node of six. The basis is symmetric; basis function $k \in \{0, \ldots, n_p - 1\}$ is the mirror image of basis function $n_p - k - 1$. Thus, the blue and cyan functions are mirror images around reference coordinate 0.

(779–782) In addition, the basis is symmetric, meaning basis function $\phi_i^{n_p}(x) = \phi_{n_p-1-i}^{n_p}(-x)$. Thus, first, support nodes are specified for regions 0 through $\lfloor n_p/2 \rfloor - 1$, and the support nodes for the remaining regions are determined by symmetry. Second, if n_p is even, then the middle region, $r = n_p/2 - 1$, has support nodes $\mathcal{I}_r^{n_p}$ that are symmetric around reference coordinate 0.

RC1: *l195: "L provides a basis for degree-d polynomials." I would say that the basis functions are the Π_i functions defined in the equation of L; L itself is the interpolant function defined by the basis and the specific nodal values $y(i)$.*

Thanks. This material will appear in a companion paper.

RC1: *l195: "These are supported by $n = d + 1$ points, each an element in the n-vector xn." To be consistent with the literature I would use the term "node" instead of "point".*

In the revised manuscript, I use "node" when referring specifically to GLL nodes but continue to use "point" when referring to general grid points.

RC1: *l204: "Given a departure point x" These properties define the interpolant functions, thus there's no need to say that x is a departure point, it can be any point within the element.*

Thanks. This material will appear in a companion paper.

RC1: *l208: I think this constraint is equivalent to saying that the basis must Lagrangian or nodal?*

This material will appear in a companion paper. The basis need not be Lagrangian (e.g., it turns out to be piecewise polynomial), but it must admit a nodal representation.

RC1: *l246: only here you define a d-degree polynomial. This would be useful already in section 2.*

This material will appear in a companion paper.

RC1: *l257: what is Runge's phenomenon? what is Lebesque constant? help the reader to understand the rationale behind your work.*

This material will appear in a companion paper.

RC1: *Section 3.6: the description of the search algorithm is quite technical and could perhaps be moved to the appendix; it is not necessary to follow the main storyline of the paper.*

This material will appear in a companion paper.

RC1: *Section 4: mention TTPR already in the introduction as it seems to be relevant for the entire Islet method.*

Thanks. I have moved the definition to Sect. 1.3:

(104–106) We refer to this approach as *tracer transport p-refinement* (TTPR). In the finite element method, p-refinement means increasing the basis polynomial degree. In the Islet method, we increase n_p^{t} relative to n_p^{v} to represent the mixing ratio fields at higher resolution.

RC1: *l439: earlier tracer tendencies were denoted by $f_i \Delta t$*

Thanks. I have now written the defintion explicitly:

(410–412) Second, $\Delta\mathbf{q} \equiv \mathbf{f}\Delta t$ is remapped from the physics grid to the tracer grid, where Δt is the physics parameterization time step. Either of $\mathbf{f}$ or $\Delta\mathbf{q}$ is sometimes called a *tendency*.

RC1: *l442: "immediate element neighbors" Are these neighbors that share an edge or vertex?*

Thanks. I now write:

(426–427) Let an element's neighborhood contain itself and every other element that shares a vertex with it.

RC1: *l464: "In contrast, ..." Meaning unclear, please revise.*

Thanks. I have rewritten the DSS discussion and moved it to its own subsection, Sect. 2.4. Section 2 builds up notation and concepts step by step, so the new explanation of the DSS and the generalized DSS hopefully is clearer.

RC1: *l507: What is "tracer density"? Is it just ρ? Then, for clarity, I'd call it "density"*

ρ is the air density, where "air" is the total content of the parcel. The tracer density is ρq, where q is the tracer mixing ratio. In a previous comment, I addressed how I have clarified these concepts in the text.

RC1: *section 4.3: while computational efficiency is important, I would move this section to the appendix, as it*

This material will appear in a companion paper.

RC1: *l570: "is proportional to the number of grid points" Should read: "proportional to the square of number of grid points"*

This material will appear in a companion paper. However, by "grid points" I mean general points in a grid, regardless of spatial dimension. In Sect. 2, I explain the tensor-product grid in greater detail, hopefully clarifying that essentially none of the new manuscript is concerned with 1D grids.

RC1: *l572: what is a naive h-halo exchange and how does it differ from what is proposed here?*

This material will appear in a companion paper, and I will explain the concepts in greater detail. "Naive" means deterministically exchanging data in all elements that *might* have a source-target relationship. Instead, in my implementation, I exchange only the data that are needed to determine the solution based on the flow-dependent domain of dependence in each time step, substantially reducing the communication volume. Thus, in general, the structure of the communicated data changes in each communication round.

RC1: *section5, l606: please mention the problem domain (full sphere?) and the equations that are being solved (pure advection on the tracer grid?)*

Thanks. I now write:

(502–503) Except in Sect. 4.3, the equation is the sourceless advection equation, Eq. 4. Two-dimensional, time-dependent flow, $\mathbf{u}(\mathbf{x}, t)$, is prescribed on the sphere.

RC1: *l631: why can you not use the dynamical core to compute the density? the chosen approach seems rather ad-hoc; can you guarantee that density values are realistic?*

Thanks. The new Appendix B derives the ISL discretization of the continuity equation, starting with the Reynolds transport theorem and then describing the details of the numerical quadrature, step by step, that lead to the discretization.

The experiments are run with a standalone program that is not part of the dynamical core. In addition, to run many long experiments, which are used in stability verification, I need a fast method to compute density; thus, an interpolation semi-Lagrangian method makes sense. Note

that the air density is used for only property preservation calculations, so the details of its solution are not impactful.

RC1: *scaling figures 6-8, 11-16: for easier readability do not use red line twice for np=6 and np=12. x axis label is missing.*

Thanks. I have added the label and used a different color for each value of n_p^t.

RC1: *l703: "has even more accuracy" Can you quantify this? Is the difference significant?*

Thanks. I have added a number of quantitative comparisons, as follows, where the first passage addresses the particular line in your question.

(590–599) For example, the most accurate shape-preserving method in TR14 for the nondivergent flow with Gaussian hills IC is HEL-ND-CN1.0, by a substantial margin (cyan curves in Fig. 1, bottom right, of TR14). The $n_p^t = 12$ Islet scheme with the long time step, Fig. 4, is approximately three times more accurate than HEL-ND-CN1.0 in the l_2 norm at resolution 0.375° and approximately twice as accurate at resolution 3°. Yet HEL-ND is, quoting TR14, an "unphysical" method. It is run for comparison with the practically useful HEL scheme. After HEL-ND, the next most accurate method in the l_2 norm at 0.375° resolution is CSLAM-CN5.0. The Islet method with the long time step, the same as that of CSLAM-CN5.0, is at least as accurate for $n_p^t \geq 8$. With the short time step, the same as that of CSLAM-CN1.0, the Islet method is at least as accurate as CSLAM-CN1.0 for $n_p^t \geq 6$. At 3° resolution, no method other than HEL-ND-CN1.0 provides l_2 norm below 10^{-2}; the Islet method does for $n_p^t \geq 8$ with the long time step and $n_p^t \geq 10$ (only $n_p^t = 12$ is shown) with the short time step.

(609–610) For example, with a long time step, for $n_p^t = 8$, this value is a little coarser than 3°; for $n_p^t = 12$, approximately 6°. For comparison, no model in TR14 reports a value larger than 2.5°.

(627–628) There is no summary number that can be compared directly with the results in Fig. 5 of TR14, but, visually, the curves for $n_p^t \geq 8$, resolution 1.5°, and on the tracer grid are among the best of those in TR14.

(642–644) In Figs. 11–14 in TR14, the smallest value of l_r at 1.5° among the property-preserving methods is 2.15×10^{-4}, by the UCISOM-CN5.5 method, except for a value of 0 by HEL-ND, which, again, cannot be used in practice. For the long time step, the Islet method gives at least as small a value for $n_p^t \geq 6$; for the short time step, $n_p^t \geq 8$.

(648–656) This diagnostic is more difficult to compare than l_r because very dissipative methods tend to have a large value of l_r and consequently a very small value for l_u. In contrast, a very accurate method, for which l_r is small, can have a larger l_u value than a very dissipative method. One means of comparison is to consider the best l_u values among methods that obtain, say, $l_r \leq 5 \times 10^{-4}$. In Figs. 11–14 in TR14, the smallest value of l_u at 1.5° under this restriction is 0, obtained by the HEL-CN1.0 and HEL-CN5.5 methods. These HEL variants are practically usable, unlike HEL-ND, and are designed to preserve tracer correlations exactly. Other than the HEL methods, the next best value is 4.80×10^{-5}, again by the UCISOM-CN5.5 method. For both the long and short time steps, at 1.5°, the Islet method gives at least as small a value for $n_p^t \geq 8$. However, even with the constraint on l_r, comparison is not straightforward, as the UCISOM methods are not strictly shape-preserving and so have $l_o > 0$.

(653–655) Although TR14 does not provide error norm values for this problem, those in Fig. 7 of TS12 can be compared with the Islet method's values at 1.5° resolution and the long time step (left column of Fig. 14). The Islet method's values of l_2, l_{inf}, ϕ_{min}, and ϕ_{max} are at least as good as those in Fig. 7 of TS12 for $n_p^t \geq 6$.

RC1: *l711: "For each value ..." unclear sentence, please revise.*

I have revised this passage as follows:

(618–621) For each possible value τ of the tracer mixing ratio at the initial time, the area over which the mixing ratio is at least τ at the midpoint time is computed. For the cosine bells IC, $\tau \in [0.1, 1]$. The diagnostic is then this area divided by the correct area, which for nondivergent flow is the area at the initial time. The perfect diagnostic value is 100% for all $\tau \in [0.1, 1]$ and 0 otherwise.

RC1: *l741: 0 in m(0) stands for time t=0?*

Thanks. Yes. I have now made clear the dependence on space and time:

(670–674) Two tracer mixing ratios, a source $q_1 = s$ and a manufactured tracer $q_2 = m$, are paired. At time $t = 0$, $m(\mathbf{x}, t)$ is set to 0, where $\mathbf{x}$ is position on the sphere. A tendency Δm is applied to m on the physics grid: $\Delta m(\mathbf{x}, t) \equiv -[\cos(2\pi(t + \Delta t)/T) - \cos(2\pi t/T)]s(\mathbf{x}, t)/2$, so that the exact solution is $m(\mathbf{x}, t) = (1 - \cos(2\pi t/T))s(\mathbf{x}, t)/2$ and, in particular, $m(\mathbf{x}, T/2) = s(\mathbf{x}, T/2)$.

RC1: *l780: what is a terminator?*

I have added the definition to the text:

(684–686) The reactions are extremely sensitive to solar insolation. The sun's position is held fixed with respect to the grid. As a result, the largest-scale spatial pattern one sees in the fields is the boundary dividing nonzero (day) and zero (night) solar insolation, the solar *terminator*; this boundary is particularly visible in the right image of Fig. 17...

RC1: *l897: Here you define the basis functions ϕ_i. This notation would be useful throughout the manuscript, already in Sect 2 and 3.1.*

Thanks. I now introduce the basis functions with the symbol ϕ in Sect. 2.

RC1: *abstract and intro: you define an abbreviation "dycore". I would omit it as it does not really save space, and it is not used frequently in the paper.*

Thanks; I have.

RC1: *l745: "shows the results"*

Thanks; fixed.

RC1: *l815: SYPD numbers printed above ...*

Thanks; fixed.

RC2: *This manuscript describes an interpolation-based semi-Lagrangian (SL) method for the transport problem on spectral-element (SE) domains. The SL transport schemes are widely used for multi-tracer transport in atmospheric models due to their accuracy and computational efficiency. The classical SL method employs interpolation at the upstream locations of the backward trajectories to estimate the advecting scalar values at the new time level. However, such an approach is not conservative per se, for practical applications an arbitrary procedure known as the "mass-fixing" usually employed for global conservation — which may have an adverse effect for climate-scale (long term) integration due to the local mass drifting. On the other hand, a finite-volume formulation of the SL method is conservative by design, where the upstream interpolation over the Lagrangian element is replaced by integration constrained to be locally (hence globally) mass conservative. The conservative data transfer from regular Eulerian grid to the deformed Lagrangian grid often referred to as the remapping (re-zoning), a limiter or shape-preserving scheme is usually employed for physically realizable solutions. A wide body of literature is available for both conservative and classical SL methods.*

Re: "'mass-fixing' ... may have an adverse effect for climate-scale (long term) integration due to the local mass drifting": I have attempted to clarify in a number of spots the relationship between local mass conservation and computational efficiency. In the introduction, I write:

(61–65) Another quality of a transport method is important: its computational efficiency. Computational efficiency is some measure of solution accuracy for a given set of computational resources.

Thus, when developing a new transport method, the objective is to obtain high efficiency, as measured by diagnostic values and computational cost, constrained by the need to couple to specific dynamics and physics grids. Our objective in this article is to extend our highly efficient tracer transport method for EAM version 2...

Then, later:

(173–175) Our objective in this work is to use the freedom provided by giving up local, but not global, mass conservation to maximize computational efficiency.

To address the viability of our method in long time integrations, I cite a model that uses one variant of it:

(22–24) For EAM version 2, we developed a new tracer transport method that is 6.5 to over 8 times faster than in EAM version 1, in the cases of, respectively, low and high workload per computer node [4] (Fig. 3).

RC2: *Implementation of conservative SL method on spherical domains tiled with high-order spectral-elements are very challenging. Authors have proposed an interpolation-based SL method Islet for the SE discretization. Instead of using the unstable native high-order interpolator (basis function) they have devised a cumbersome numerical procedure which employs an alternative grid system within each spectral element, adding another layer of complexity. The Islet method is not conservative, nevertheless, the global conservation is achieved by mass-fixing. The authors argue that the Islet method can handle tracer transport as well as the remapping between physics & dynamics grids, and incorporate shape-preservation filters.*

While the method is complex, EAM version 2 successfully uses two grids, incorporating the new physics grid described in this paper and [5]. I have added the following text to support this point, where the text refers to EAM version 2:

(24–26) In addition, we developed remap operators to remap data between separate grids for physics parameterizations and dynamics, permitting the physics parameterization computations to run on a coarser grid and thus 1.6 to 2.2 times faster in version 2 than in version 1 [5, 4].

To clarify that the new basis functions are not any more difficult to use than the natural basis functions, I added this text to the first paragraph describing the Islet basis functions in Appendix A:

(764–765) This article can be understood equally well by assuming the standard GLL bases are used; only the numerical results depend on the details of the basis functions.

RC2: *The manuscript is very long, the Islet interpolation as described by the authors is extremely complex. Authors failed to explain the core interpolation algorithm with clarity, there are many statements in the manuscript which leads to ambiguity. The numerical analysis part is very intense maybe more suitable for a computational math journal (e.g., SIAM / JCP) than the GMD. The subject covered could be split into a two-part paper, one describing the basic algorithm and analysis with more details and rigor, and the second part for implementation and validation with standard tests. This would be helpful for better reading. The current manuscript is written in an awkward manner and is unacceptable for publication.*

Thanks. I have split the original manuscript into two. To improve clarity, I wrote the new Sect. 2. Section 2 describes each algorithmic subcomponent, with the following outline corresponding to subsections:

- Semi-Lagrangian transport
  - Types of SL methods
  - Spectral element ISL transport

- Definitions and notation
- The linear advection operator

- Property preservation
  - ClipAndAssuredSum
  - Infeasible problems
  - Local and global problems
  - Relaxed problems

- Grid remap

- Direct stiffness summation

Each subsection provides a text overview followed by mathematical details. Section 3 then directly references each subsection from Sect. 2 as it discusses each step of the Islet method.

RC2: *Recommendation: Major revision, possibly resubmit as a two-part manuscript. Authors should address the following questions.*

I have separated the manuscript into two.

RC2: *(1) The stability associated with the SL method is that the deformational Courant number (Lipschitz condition) should not exceed unity, in plain language, the trajectories should not cross intersect (see, Staniforth & Cotes 1992 MWR paper). Is the cubic ISL method (lines 115-120) unstable due to this condition? Need some explanation.*

I addressed this comment in my previous response to reviewer 2. This material is used in another manuscript.

RC2: *(2) The SL transport scheme can be stabilized using a limiter, filter or with an explicit diffusion (see, Ullrich & Norman, QJRMS, 2014). You can use the native high-order SE interpolation (basis function) for the SL transport combined with the limiter which you are already using for the Islet method. It will be interesting to see how the Islet method compares with this simple SL-SE scheme employing 4x4 GLL grid (I guess that is the SE grid choice made for the operational E3SM).*

I addressed this comment in my previous response to reviewer 2. Since the material on stability now is used in a second manuscript, I have moved the figure and text demonstrating the effect of instability on the method using "the native high-order SE interpolation (basis function) for the SL transport combined with the limiter" to Appendix A3.

RC2: *(3) It is not convincing to have 3 grid systems (physics: FV, dynamics: GLL, transport: tweaked GLL) in a SE modeling framework. The Fig.5 shows such a grid configuration, and it appears to be very challenging. At a very high (NH) resolution the data movement is a major issue for an element-based Galerkin model (DG/SE). A typical climate model may have O(100) tracers, an additional tracer grid with more DOF than the dynamic grid can exacerbate this problem. This will limit the use of Islet scheme, how do you address it?*

Thanks. I addressed this comment in my previous response to reviewer 2. I have also added the following text to the introduction, where Fig. 3 in [4] is Fig. 1 in my response to reviewer 2 (https://doi.org/10.5194/gmd-2021-296-AC3):

(18–26) Because of the large number of tracers in climate models, tracer transport can be computationally very expensive. For example, in the Dept. of Energy's Energy Exascale Earth System Model (E3SM) [2] Atmosphere Model (EAM) version 1 [3], configured with the default 40 tracers, tracer transport takes approximately 75% of the total dynamical core wall clock time and

approximately 23% of the total atmosphere model wall clock time on a typical computer cluster [4] (Fig. 3). For EAM version 2, we developed a new tracer transport method that is 6.5 to over 8 times faster than in EAM version 1, in the cases of, respectively, low and high workload per computer node [4] (Fig. 3). In addition, we developed remap operators to remap data between separate grids for physics parameterizations and dynamics, permitting the physics parameterization computations to run on a coarser grid and thus 1.6 to 2.2 times faster in version 2 than in version 1 [5, 4].

RC2: *(4) With real data you have velocity information only available at the GLL (dynamics) grid, the way you find the 2D trajectory information using the 3D Cartesian coordinates leads to additional computational overhead when the method is extended to the 3D application (line 470-475). This needs some justification, why not use the spherical (u,v) components or corresponding contravariant vectors?*

I addressed this comment in my previous response to reviewer 2.

RC2: *It is not clear that the maximum eigenvalue required for the interpolation is the tracer data dependent, in that case you have a serious computational overhead for the multi-tracer applications, Please clarify! What is the computational halo requirement for an SE stencil with NxN GLL points, when the shape preserving limiter is applied?*

I addressed this comment in my previous response to reviewer 2.

RC2: *(5) What is the special advantage of using Islet method? It seems you have introduced a complex numerical method for a relatively simple linear transport problem. If mass-fixing is the way to go, one could use the RBF-based (Kriging type) interpolator which provides very accurate solution, and no need for the expensive search for max eigenvalue etc.*

I addressed this comment in my previous response to reviewer 2.

RC2: *(6) The results are looking good, authors should limit the number of figures and make an effort to compare the results with that of other high-order element-based schemes. Why the results from your own previous papers (Bosler et. al. 2019, SIAM J Sci. Computing; Guba et al. 2014, JCP) discussed? These results should be compared and the relative merits should be discussed.*

Thanks. I have added a number of comparisons to the text, as documented above in a response to reviewer 1.

Re: Guba et al. 2014, JCP, results from that method appear in TR14 and so are included in the above comparisons. In addition, the performance of our ISL method is compared in lines 18–26 of the manuscript quoted above.

Re: Bosler et. al. 2019, SIAM J Sci. Computing, I have added the following text to make clear why we are pursuing an *interpolation* method instead of an *exactly cell-integrated* method:

(152–171) Interpolation is in contrast to *exactly cell-integrated* methods, which accurately integrate the basis of a target (e.g. Lagrangian) element against those of the source; see, e.g., [1]. (In some cases, an inaccurate cell-integrated method can be interpreted as an interpolation method; see Appendix B for an example.) Exactly cell-integrated methods have substantially greater cost than interpolation methods for three reasons.

First, to obtain smoothness in the integrand, integration is over facets computed by geometric intersection of a target element against source elements; intersection calculations are not needed in interpolation methods. Typically, to minimize computational geometry complexity, departure cell edges are approximated by great arcs rather than flow-distorted curves, limiting the method to second-order accuracy; however, [6] describe a higher-order edge reconstruction that yields a third-order accurate advection method. In contrast, achieving arbitrarily high order in an ISL method's linear advection operator does not entail any additional complexity.

Second, accurate integration has a larger computational cost because it requires sphere-to-reference point calculation and interpolant evaluations at many quadrature points.

Third, an exactly cell-integrated method requires a larger communication volume because all data from a source element are used to integrate against each target basis function.

In trade for these additional costs, exactly cell-integrated methods are locally mass conserving, and the fact that they are L^2 projections can be used to prove stability. Local mass conservation means that one can identify numerical, possibly Lagrangian, fluid parcels on the grid that have constant tracer mass. *Local* is in contrast to *global* mass conservation; the latter means that the mass of the tracer fluid is conserved over the whole domain but not necessarily in any identifiable parcels smaller than the domain. Although an exactly cell-integrated method is locally mass conserving, coupling it to a dynamics solver still generally requires additional measures to obtain mass-tracer consistency.

RC2: *(7) There are many undefined terms (e.g. CAAS) and notations which I am going to list, this should be fixed.*

Thanks. CAAS in particular is now explained in detail in Sect. 2.2.1. I have attempted to make the notation very clear in Sect. 2 to make Sect. 3 easier to read.

Thanks,
Andrew Bradley

References

[1] P. A. Bosler, A. M. Bradley, and M. A. Taylor. Conservative multimoment transport along characteristics for discontinuous galerkin methods. *SIAM J. on Sci. Comput.*, 41(4):B870–B902, 2019.

[2] E3SM Project. Energy Exascale Earth System Model (E3SM). [Computer Software] https://dx.doi.org/10.11578/E3SM/dc.20180418.36, April 2018.

[3] J.-C. Golaz, P. M. Caldwell, L. P. Van Roekel, M. R. Petersen, Q. Tang, J. D. Wolfe, G. Abeshu, V. Anantharaj, X. S. Asay-Davis, D. C. Bader, et al. The DOE E3SM coupled model version 1: Overview and evaluation at standard resolution. *J. Adv. Model Earth Sy.*, 11(7):2089–2129, 2019.

[4] J.-C. Golaz et al. The doe e3sm model version 2: Overview of the physical model. *in preparation for J. Adv. Model Earth Sy.*, 2022.

[5] W. M. Hannah, A. M. Bradley, O. Guba, Tang Q., J.-C. Golaz, and J. Wolfe. Separating physics and dynamics grids for improved computational efficiency in spectral element earth system models. *J. Adv. Model Earth Sy.*, 13(7):e2020MS002419, 2021.

[6] P. A. Ullrich, P. H. Lauritzen, and C. Jablonowski. Some considerations for high-order incremental remap-based transport schemes: edges, reconstructions, and area integration. *Int. J. for Numer. Meth. in Fl.*, 71(9):1131–1151, 2013.