Mathematics Faculty Publications and Presentations

Department of Mathematics

Mesh-Free Semi-Lagrangian Methods for Transport on a Sphere Using Radial Basis Functions

Varun Shankar
University of Utah

Grady B. Wright
Boise State University

Mesh-free Semi-Lagrangian Methods for Transport on a Sphere Using Radial Basis Functions

Varun Shankar
Dept. of Mathematics
University of Utah
Salt Lake City, UT 84112

Grady B. Wright
Dept. of Mathematics
Boise State University
Boise, ID 83725

Abstract

We present three new semi-Lagrangian methods based on radial basis function (RBF) interpolation for numerically simulating transport on a sphere. The methods are mesh-free and are formulated entirely in Cartesian coordinates, thus avoiding any irregular clustering of nodes at artificial boundaries on the sphere and naturally bypassing any apparent artificial singularities associated with surface-based coordinate systems. For problems involving tracer transport in a given velocity field, the semi-Lagrangian framework allows these new methods to avoid the use of any stabilization terms (such as hyperviscosity) during time-integration, thus reducing the number of parameters that have to be tuned. The three new methods are based on interpolation using 1) global RBFs, 2) local RBF stencils, and 3) RBF partition of unity. For the latter two of these methods, we find that it is crucial to include some low degree spherical harmonics in the interpolants. Standard test cases consisting of solid body rotation and deformational flow are used to compare and contrast the methods in terms of their accuracy, efficiency, conservation properties, and dissipation/dispersion errors. For global RBFs, spectral spatial convergence is observed for smooth solutions on quasi-uniform nodes, while high-order accuracy is observed for the local RBF stencil and partition of unity approaches.

Key words: RBF, Hyperbolic PDEs, Advection, meshless

1 Introduction

Radial basis function (RBFs) methods have been used for over a decade to solve partial differential equations (PDEs) on spheres. These methods can broadly be classified into *global* RBF collocation methods [1, 2, 3, 4, 5], RBF-generated finite difference (RBF-FD) methods [6, 7, 8, 9, 10], and more recently RBF-partition of unity (RBF-PU) collocation methods [11]. Global RBF methods when used with infinitely-smooth RBFs show spectral convergence on smooth problems at the cost of *dense* differentiation matrices; in contrast, RBF-FD and RBF-PU methods produce sparse differentiation matrices and high-order algebraic convergence rates. All of these methods can use "scattered" nodes in their discretizations of a sphere, and have the benefit of being independent of any surface-based coordinate system. They thus avoid any unnatural grid clustering and do not suffer from any coordinate singularities.

In this paper, we present three new RBF methods with similar benefits for numerically solving the transport equation on the surface of a sphere in an incompressible velocity field. For the unit sphere $\mathbb{S}^2$, this PDE is given

$$\frac{Dq}{Dt} = 0, \quad \frac{D}{Dt} := \frac{\partial}{\partial t} + \boldsymbol{u} \cdot \nabla_{\mathbb{S}^2}, \tag{1}$$

where q is the scalar quantity being transported, $\boldsymbol{u}$ is a surface divergence-free vector field that is tangent to $\mathbb{S}^2$, and $\nabla_{\mathbb{S}^2}$ denotes the surface gradient operator on $\mathbb{S}^2$. Since global atmospheric flows are dominated by the horizontal advection process, the numerical solution to the transport problem is a fundamental part of any solver for these flows.

Currently, all RBF discretizations for the transport equation (and more general hyperbolic equations like the shallow water equations) on a sphere suffer from the same drawback: the eigenvalues of the differentiation

matrices corresponding to the surface gradient operator may in general have positive real parts, leading to either a slow or rapid onset of instability during a numerical simulation [1]. As of this paper, the only known approach to rectify this problem is to add an artificial hyperviscosity of the form $(-1)^{(k+1)}\gamma\Delta^k$ in the right hand side of the PDE, where $k \geq 2$ is an integer and $\gamma > 0$ is some small real number that scales *inversely* with the total number of nodes N. The intuition here is that higher powers of the Laplacian Δ will damp out the eigenvectors associated with the rogue eigenvalues of the discretized surface gradient, while leaving the others essentially untouched [6, 7, 12, 13]. With global RBFs, the hyperviscosity operator typically takes the form of γA^{-1}, where A is the global RBF interpolation matrix whose inverse mimics the properties of high powers of the Laplacian [6]; a similar approach can be employed for the RBF-PU method [11]. Unfortunately, for a given PDE and node set on a sphere, the precise values of γ and k required to stabilize the numerical solution may need to be determined by trial and error, which can add to the computational expense of the method.

A common way to naturally stabilize local Eulerian methods for transport is to use "upwinding", which uses dynamic direction-dependent information about the flow field. However, this form of upwinding typically requires an underlying mesh and so is impractical for truly mesh-free local methods like RBF-FD and RBF-PU collocation. An important class of methods that naturally possess upwinding are semi-Lagrangian (SL) schemes, widely acknowledged for their formal independence from the CFL stability condition [14]. SL methods have successfully been used for simulating various problems in fluid dynamics, especially in numerical weather prediction and climate modeling, from simple tracer transport to more complex problems involving wide ranges of spatiotemporal scales and intricate forcing terms, e.g. [14, 15, 16, 17, 18]. For problems on a sphere, SL methods have generally used latitude-longitude grids and spherical coordinate systems (e.g. [19, 20, 21]), or other regular surface-based grids and local surface-based coordinate systems (e.g. [22, 23, 24]), which can lead to a loss of accuracy because of singularities that arise in the mappings from the physical sphere to the surface-based coordinate systems.

In this paper, we present three new high-order SL methods for transport on a sphere based on interpolation with global RBFs, local RBFs, and RBF-PU methods. These methods are mesh-free, allowing for scattered node discretizations, and are formulated entirely in Cartesian coordinates so as to avoid any surface based coordinate singularities. The local RBF and RBF-PU methods also allow for a type of "*p*-refinement" for increasing the accuracy for a given fixed set of discretization nodes. We demonstrate that the SL framework lends our new methods both accuracy and intrinsic stability, thereby eliminating the need for a hyperviscosity term. For the local RBF and RBF-PU methods, we propose using "scale-free" RBFs appended with spherical harmonics (an idea related to that of Flyer et al. [25] for planar domains) to further reduce the number of tuning parameters (*i.e.*, the shape-parameter) and to bypass so-called error stagnation. We compare and contrast all three methods using three standard test cases from the literature—solid-body rotation of a cosine bell from [26] and deformational flow of two bells from [27]. The focus of these comparisons is on the overall accuracy, dissipation and dispersion properties, mass conservation, and computational cost. We find that the computational costs of these new methods are comparable to those of existing RBF collocation and finite-difference techniques. In particular, we find that the local RBF and RBF-PU methods are highly scalable to large node sets. We note that RBFs have previously been used for SL advection in [28], but the focus there was on planar domains and global RBF methods. Additionally, they have been used in a conservative SL advection method in [29] for planar domains using Voronoi cells and local thin plate splines. This is the first application of RBFs to SL transport on a sphere with accurate and scalable numerical methods and in a entirely mesh-free formulation.

We note that, while SL methods are applicable to a wide class of advection dominated problems, they have some limitations, the primary one being a lack of local conservation unless explicitly formulated in a conservative manner. However, the large-time-step conservative SL methods are computationally expensive [30].

The paper is organized as follows. In the next section, we review the global RBF, local RBF and RBF-PU methods in the context of interpolation. In Section 3, we review the SL advection technique and discuss how to use use the three RBF methods within the SL framework in an efficient fashion. In Section 4, we compare and contrast the new SL methods on three standard test cases for transport on the unit sphere. Finally, we conclude with a summary of our results and future research directions in Section 5.

2 Global, local, and partition of unity RBF interpolation on $\mathbb{S}^2$

RBFs are a well-established method for interpolating/approximating data over a set of "scattered" nodes $X = \{\boldsymbol{x}_j\}_{j=1}^N \subset \Omega \subseteq \mathbb{R}^d$. The standard method uses linear combinations of shifts of a kernel $\phi : \Omega \times \Omega \to \mathbb{R}$ with the property that $\phi(\boldsymbol{x}, \boldsymbol{y}) := \phi(\|\boldsymbol{x} - \boldsymbol{y}\|)$ for $\boldsymbol{x}, \boldsymbol{y} \in \Omega$, where $\| \cdot \|$ is the Euclidean norm in $\mathbb{R}^d$. Kernels with this property are referred to as *radial kernels* or simply *radial functions*. In the case where $\Omega = \mathbb{S}^2$, these kernels are sometimes referred to as *spherical basis functions* since $\phi(\|\boldsymbol{x} - \boldsymbol{y}\|) = \phi(2 - 2\boldsymbol{x}^T\boldsymbol{y})$ when $\boldsymbol{x}, \boldsymbol{y} \in \mathbb{S}^2$, i.e. ϕ will only depend on the cosine of the angle between $\boldsymbol{x}$ and $\boldsymbol{y}$. We will use the term RBFs and not make use of this simplification since our numerical schemes use extrinsic coordinates.

Below we review the three interpolation methods used in this study. We give cursory details on the first two methods, as they have appeared in many other places in the literature; see the recent book by Fornberg & Flyer [31] for further details and applications of these two methods.

2.1 Global RBF interpolation

Given a set of nodes $X = \{\boldsymbol{x}_k\}_{k=1}^N \subset \mathbb{S}^2$ and a continuous target function $f : \mathbb{S}^2 \to \mathbb{R}$ sampled at the nodes in X, the standard *global* RBF interpolant to the data has the form

$$s(\boldsymbol{x}) = \sum_{k=1}^N c_k \phi(\|\boldsymbol{x} - \boldsymbol{x}_k\|). \tag{2}$$

The expansion coefficients $\{c_k\}_{k=1}^N$ are determined by enforcing $s|_X = f|_X$, which can be expressed by the following linear system:

$$\underbrace{\begin{bmatrix} \phi(r_{1,1}) & \phi(r_{1,2}) & \dots & \phi(r_{1,N}) \\ \phi(r_{2,1}) & \phi(r_{2,2}) & \dots & \phi(r_{2,N}) \\ \vdots & \vdots & \ddots & \vdots \\ \phi(r_{N,1}) & \phi(r_{N,2}) & \dots & \phi(r_{N,N}) \end{bmatrix}}_{A_X} \underbrace{\begin{bmatrix} c_1 \\ c_2 \\ \vdots \\ c_N \end{bmatrix}}_{\mathbf{c}_X} = \underbrace{\begin{bmatrix} f_1 \\ f_2 \\ \vdots \\ f_N \end{bmatrix}}_{\mathbf{f}_X}, \tag{3}$$

where $r_{i,j} = \|\boldsymbol{x}_i - \boldsymbol{x}_j\|$. If ϕ is, for example, a positive-definite radial kernel on $\mathbb{R}^3$, and all nodes in X are distinct, then the matrix A_X above is guaranteed to be positive definite, so that (2) is well-posed. Examples of various choices for ϕ, including relaxed conditions to guarantee the well-posedness of (2) can be found in [32, Ch. 3–12]. Results on the approximation properties of RBF interpolants on the sphere for target functions of various smoothness can be found in, for example, [33, 34, 35]. A result from [35] is that for infinitely-smooth target functions, convergence rates that are faster than any polynomial order can be realized for various infinitely-smooth ϕ. Additionally, global RBF interpolants based on various infinitely-smooth kernels have been shown to converge to spherical harmonic interpolants as the kernels are scaled to become flat (the so-called flat limit) [36].

One issue with global RBF interpolation is that the linear system for determining the interpolation coefficients (3) is dense, thereby leading to an $O(N^3)$ computational cost to solve it using a direct method. Another issue is that the matrices can become ill-conditioned. While the RBF-QR method [36] can be used to bypass the ill-conditioning associated with small shape parameters (i.e. flat radial kernels), algorithms for bypassing the $O(N^3)$ computational cost *and* ill-conditioning–without sacrificing high-order accuracy–have not yet been developed. The next two subsections discuss techniques for addressing these issues. We note that an alternative to these methods, that is not pursued here, is to use compactly-supported kernels in a multilevel type framework [37].

2.2 Local RBF interpolation

In this method, RBF interpolants are used locally over a small collection of points surrounding each node in the set X (this technique forms the foundation for RBF-FD methods [38]). The method proceeds as follows. For each node $\boldsymbol{x}_k$, $k = 1, \dots, N$, in the set X, we select subsets of X that consist of $\boldsymbol{x}_k$ and its

$n - 1$ nearest neighbors, where $n << N$. We refer to these subsets as *stencils* and denote them by X_k, and refer to node $\boldsymbol{x}_k$ as the *center point* of the stencil X_k; see Figure 1 for an illustration. The nearest neighbors are typically determined in a preprocessing step using a data structure such as a kd-tree, which typically requires $O(N \log N)$ operations to construct. On each stencil, an RBF interpolant is constructed to the nodes in X_k and the corresponding function values. It is then used to reconstruct the underlying target function globally at all points within some neighborhood of the center point of the stencil. While this approach does not produce a globally smooth interpolant to the underlying target function, we find that it works well in the SL setting.


Figure 1: Illustration of the nodes in an example stencil for the local RBF method on the sphere. The global nodes X are marked as small greyed balls, while the nodes that make up the stencil X_k are marked as black balls, with the center point $\boldsymbol{x}_k$ marked in red an with a black square around it.

For the local RBF interpolants, rather than use (2) on each stencil, we follow an approach related to that proposed by Flyer et al. [25] for RBF-FD methods to avoid stagnation errors for interpolation in $\mathbb{R}^d$. Their approach augments the standard form of the interpolant (2) with a linear combination of polynomials up to a degree corresponding to roughly $n/2$ polynomial basis functions, together with a set of "moment conditions" on the RBF coefficients. A natural extension of their approach to the sphere is to augment the standard interpolant with *spherical harmonics* up to a degree L that grows with n (see Section 2.4), since these are polynomials in $\mathbb{R}^3$ that are linearly independent when restricted the sphere. The exact form of the augmented local RBF interpolants on each stencil X_k that we use is then given by

$$s_k(\boldsymbol{x}) = \sum_{j=\mathcal{I}_1^k}^{\mathcal{I}_n^k} c_j^k \phi(\|\boldsymbol{x} - \boldsymbol{x}_j\|) + \sum_{i=1}^{(L+1)^2} d_i^k p_i(\boldsymbol{x}), \tag{4}$$

where $\mathcal{I}_1^k, \ldots, \mathcal{I}_n^k$ are the indices into the global node set X of the nodes contained X_k and $\{p_i\}_{i=1}^{(L+1)^2}$ are a basis for the space of spherical harmonics of degree L. To determine the unknown coefficients, we impose that s_k interpolates the data associated with the stencil X_k and that the coefficients c_j^k satisfy the following $(L+1)^2$ moment conditions:

$$\sum_{j=\mathcal{I}_1^k}^{\mathcal{I}_n^k} c_j^k p_i(\boldsymbol{x}_j) = 0, \quad i = 1, \ldots, (L+1)^2,$$

which are the standard conditions imposed in the literature for polynomials in $\mathbb{R}^d$ [25, 32]. The interpolation and moment conditions lead to the following linear system:

$$\underbrace{\begin{bmatrix} A_k & P_k \\ P_k^T & 0 \end{bmatrix}}_{\widetilde{A}_k} \underbrace{\begin{bmatrix} \mathbf{c}_k \\ \mathbf{d}_k \end{bmatrix}}_{\widetilde{\mathbf{c}}_k} = \underbrace{\begin{bmatrix} \mathbf{f}_k \\ \mathbf{0} \end{bmatrix}}_{\widetilde{\mathbf{f}}_k}, \tag{5}$$

where A_k is the RBF interpolation matrix for the nodes in stencil X_k (see (3)) and P_k is a n-by-$(L+1)^2$ matrix with column i containing entries $p_i(\boldsymbol{x}_j)$, $j = \mathcal{I}_1^k, \ldots, \mathcal{I}_n^k$. These local interpolants can be used to

4

approximate functions to high-order algebraic accuracy determined by the stencil size n and the degree of the spherical harmonics. Much like in [25], we find that augmenting the standard RBF interpolant with spherical harmonics significantly enhances the accuracy of the local RBF interpolants. We will henceforth use tilde's above variables that correspond to the augmented RBF system (5).

2.3 RBF-PU interpolation

The RBF-PU method is similar to the local RBF interpolation method. However, instead of forming and using different local interpolants for each node $\boldsymbol{x}_k$, local interpolants are constructed over a collection of *patches* that cover the sphere. These interpolants are then combined using compactly-supported weight functions on each patch that, all together, form a partition of unity. In this way, the method results in a smooth interpolant over the sphere, whereas the local approach does not. The RBF-PU method was first introduced for problems in the plane by Wendland [39] (see also [32, Ch. 29]) and for interpolation problems on a sphere by Cavoretto & De Rossi in [40]. It has since extended to other domains and applications [12, 41, 11]. Specific details on the RBF-PU construction are provided below in the context of our application.

Let $\Omega_1, \ldots, \Omega_M$ be a collection of distinct spherical caps on $\mathbb{S}^2$ with the properties that the set of all patches provides an open covering of $\mathbb{S}^2$, i.e. $\cup_{\ell=1}^{M} \Omega_\ell = \mathbb{S}^2$, and each Ω_ℓ contains at least one node from X. These caps are the patches in the PU method. Figure 2 (a) illustrates a typical collection of patches for a quasi-uniformly distributed node set X on $\mathbb{S}^2$, which is what we use in this study. Let $\boldsymbol{\omega}_\ell \in \mathbb{S}^2$, $\ell = 1, \ldots, M$, denote the center of patch Ω_ℓ and $R_\ell > 0$ denote the radius of the patch, measured as the Euclidean distance from $\boldsymbol{\omega}_\ell$. For each patch, we define the following compactly-supported weight function:

$$\varphi_\ell(\boldsymbol{x}) = \varphi\left(\frac{\|\boldsymbol{x} - \boldsymbol{\omega}_\ell\|}{R_\ell}\right), \tag{6}$$

where $\varphi : \mathbb{R}^+ \to \mathbb{R}$ has compact support over the interval $[0, 1)$. In this study, we use the radially-symmetric cubic B-spline

$$\varphi(r) = \begin{cases} \frac{2}{3} + 4(r-1)r^2 & \text{if } 0 \leq r < \frac{1}{2}, \\ -\frac{4}{3}(r-1)^3 & \text{if } \frac{1}{2} \leq r < 1, \\ 0 & \text{if } 1 \leq r. \end{cases} \tag{7}$$

Using (6), we define the PU weight functions for the collection of patches $\Omega_1, \ldots, \Omega_M$ as

$$w_\ell(\boldsymbol{x}) = \frac{\varphi_\ell(\boldsymbol{x})}{\displaystyle\sum_{j=1}^{M} \varphi_j(\boldsymbol{x})}, \ \ell = 1, \ldots, M. \tag{8}$$

Note that each w_ℓ is only supported over Ω_ℓ and that the summation on the bottom only involves terms that are non-zero over patch Ω_ℓ, which should be much smaller than M. Figure 2 (b) displays one of these weight functions for a quasi-uniform distribution of overlapping patches.

We use these PU weight functions to define the RBF-PU interpolant as follows. Let X_ℓ denote the nodes from X that belong to patch Ω_ℓ and let s_ℓ denote the RBF interpolant of the form (4) to the target function f over X_ℓ. Then the RBF-PU interpolant is given by

$$s(\boldsymbol{x}) = \sum_{\ell=1}^{M} w_\ell(\boldsymbol{x})s_\ell(\boldsymbol{x}). \tag{9}$$

Since $\{w_\ell\}_{\ell=1}^{M}$ form a partition of unity, it is straightforward to show that s interpolates the target function at all nodes in X. The interpolant is also smooth over the entire sphere. Here we are using the augmented RBF interpolants (4) over each patch (which differs from [40]). We have found that this gives better accuracy than the standard interpolant (2). In Section 3.3.3, we discuss our strategy for choosing the patches.

Figure 2: (a) Illustration of a typical collection of patches (outlined in blue) for a quasi-uniformly distributed node set of $N = 4096$ nodes (marked in black) on $\mathbb{S}^2$. The centers of the patches are marked in red and there are roughly 100 nodes per patch. (b) Illustration of one of the PU weight functions (8) for the patches from part (a).

2.4 Choosing the radial kernels

In the past, infinitely-smooth radial kernels, such as the inverse multiquadric (IMQ) kernel $\phi(r) = (1 + (\varepsilon r)^2)^{-1/2}$, have been primarily used for RBF approximations to advection-dominated PDEs on a sphere, e.g. [6, 7, 1, 2, 3, 5]. When used with global RBFs, this has been demonstrated to result in convergence rates higher than any polynomial order for smooth solutions. These kernels have also been shown to be better than finitely-smooth kernels even when the solutions are not smooth [2]. However, the advantages of using infinitely-smooth kernels over finitely-smooth ones with RBF-FD was recently called into question [25, 13]. Indeed, recent work on Euclidean geometries [13, 42] has demonstrated that it is more beneficial to use the finitely-smooth polyharmonic spline (PHS) kernels in the context of RBF-FD methods, augmented with a moderate-degree polynomial. These kernels take the form $\phi(r) = r^{2k} \log r$, $k \geq 1$, or $\phi(r) = r^{2k+1}$, $k \geq 0$. The choice between these two kernels is often made depending on the dimension of the interpolation problem for theoretical reasons, while the choice of k controls the smoothness of the kernels, and hence their accuracy [32].

In this work, we use the IMQ kernel for the global RBF method in order to benefit from the higher-than-polynomial convergence rates for smooth problems. For the local RBF and RBF-PU methods we follow the recent work of [25, 13, 42, 43] and use the PHS kernel $\phi(r) = r^{2k+1}$ in conjunction with spherical harmonics, as given in (4). We choose the degree of the spherical harmonics as

$$L = \left\lfloor \frac{1}{2}(\sqrt{n} - 1) \right\rfloor,$$

and the order of the PHS as $k = L$. For the many test problems we considered, these parameters produced excellent results. Choosing $k = L$ also coincides with theoretical results on solvability of the PHS interpolation problem [32]. In contrast, recent work [25, 13, 42, 43] has used polynomials of a much higher-degree than that required to prove solvability. For problems posed in general domains in $\mathbb{R}^2$ and $\mathbb{R}^3$, it was demonstrated in these studies (see, for example, Figures 7 and 8 from [25]) that using this approach for doing local interpolations/approximating derivatives one typically obtains better accuracy. In our extensive experiments, we did not find this to carry over to the case for either the local or RBF-PU methods when the domain is $\mathbb{S}^2$ and spherical harmonics are used. In [44], the authors use a different approach for constructing local interpolants over the sphere (for doing quadrature), where the local stencil (or patch) is projected onto the tangent plane of $\mathbb{S}^2$ at the center of the stencil and the interpolations are done in the plane. This approach, which we did not test here, may give similar results for the sphere as the observations from [25, 13, 42, 43].

3 RBFs for SL advection on a sphere

As discussed in the introduction, SL advection is a technique for solving the transport equation (1). The primary idea of SL advection is to use the Lagrangian frame to find upwind directions for which q moves, and use an Eulerian set of nodes to perform interpolations of q. The upwinding step ensures that the numerical domain of dependence matches the physical domain of dependence. The presence of an Eulerian frame avoids spatial resolution issues and allows for high-order accuracy, a feature that is difficult to achieve in the Lagrangian frame without remeshing, or redistributing nodes.

The upwinding is done as follows: assume that a Lagrangian parcel arrived at each of the Eulerian nodes carrying with it some amount of the scalar q. The amount of q at this node must have been brought forward from the *departure point* of the parcel. Thus, we find the departure point by tracing the parcel along the flow field $\boldsymbol{u}$ and determine q by interpolating the scalar field to the departure point. This procedure is laid out explicitly in Algorithm 1. More details on the general SL advection procedure can be found, for example, in [45]. In the following subsections, we discuss the different aspects of the algorithm as it pertains to advection on $\mathbb{S}^2$.

Algorithm 1 Semi-Lagrangian advection on $\mathbb{S}^2$

Input: Velocity field $\boldsymbol{u}(\boldsymbol{x}, t)$ tangent to $\mathbb{S}^2$; initial scalar field $q(\boldsymbol{x}, 0)$; $X = \{\boldsymbol{x}_j\}_{j=1}^N$, $X \subset \mathbb{S}^2$; final time t_f; time-step Δt.
Set $\mathbf{q}_X^0 = \{q(\boldsymbol{x}_j, 0)\}_{j=1}^N$, $t = 0$, and $m = 0$.
while $t \leq t_f$ **do**
    For $j = 1, \ldots, N$, trace back $\boldsymbol{x}_j$ to time t to find departure point $\boldsymbol{\xi}_j^m$.
    Interpolate $\mathbf{q}_X^m$ to $\Xi^m = \{\boldsymbol{\xi}_j^m\}_{j=1}^N$ to obtain $\mathbf{q}_\Xi^m$.
    Set $\mathbf{q}_X^{m+1} = \mathbf{q}_\Xi^m$, $m = m + 1$, and $t = m\Delta t$.
end while

3.1 Eulerian node sets

Central to all SL advection schemes is a fixed set of Eulerian nodes $X = \{\boldsymbol{x}_j\}_{j=1}^N$ over the given domain that is used to interpolate the advected quantity to the Lagrangian parcels. The approximate solution to the PDE will ultimately be computed only at the set of Eulerian nodes X. Since the interpolation schemes used in this study are all based on RBFs, the Eulerian nodes are not required to live on a grid or mesh, and can be chosen freely for our application. In this study, we use node sets that provide near optimal resolution over $\mathbb{S}^2$. Since $N = 20$ nodes is the maximum number that can be exactly equally distributed over $\mathbb{S}^2$, one is resigned to using node sets that are only quasi-uniformly distributed over $\mathbb{S}^2$. These node sets, which can be generated from a variety of algorithms [46], have the property that the average spacing between nodes, h, satisfies $h \sim 1/\sqrt{N}$. In the results presented here we have used maximum determinant nodes [47] and icosahedral nodes [48]. While the latter of these sets forms a natural grid, we do not use this fact in our algorithms. All of these node sets, including many others, are available in the `SpherePts` software package [49].

3.2 Trajectory reconstruction

The process of tracing Lagrangian parcels back along the velocity field $\boldsymbol{u}$ is known as *trajectory reconstruction*. Let $\boldsymbol{\xi}_j(t) \in \mathbb{S}^2$ denote the position of a Lagrangian parcel as a function of time such that $\boldsymbol{\xi}_j(t_{m+1}) = \boldsymbol{x}_j$, for some time $t = t_{m+1} > 0$. Here $\boldsymbol{x}_j$ is the j^{th} node in Eulerian node set $X = \{\boldsymbol{x}_i\}_{i=1}^N \subset \mathbb{S}^2$. Then we find the departure point of this parcel at time $t = t_m = t_{m+1} - \Delta t$ by solving the simple ODE

$$\frac{d\boldsymbol{\xi}_j}{dt} = \boldsymbol{u}, \ \boldsymbol{\xi}_j(t_{m+1}) = \boldsymbol{x}_j, \tag{10}$$

backward in time over the interval $[t_m, t_{m+1}]$. Here Δt defines the interval over which we trace back the parcel from $\boldsymbol{x}_j$, and it also defines the time-step for solving the advection equation (1).

When the velocity field $\boldsymbol{u}$ is known for all time and space the trajectory ODE (10) can be solved in a straightforward manner using a variety of techniques. However, since we use the Cartesian coordinate representation of the velocity field $\boldsymbol{u}$ in the trajectory reconstruction, the *Lagrangian parcels* may not reside on the sphere at time $t = t_m$. Similar to [50], which also uses Cartesian coordinates (albeit in a fully Lagrangian scheme), we have found that the parcels remain very close to the sphere when a high-order integrator is used. In our experiments, we found that both the standard fourth-order Runge-Kutta (RK4) method and Fehlberg's fifth-order Runge-Kutta (RK5) scheme [51], worked well. We ultimately decided on the RK5 scheme since we found that it allowed for larger time-steps than RK4. We also found that orthogonally projecting the Lagrangian parcels exactly back to the surface of the sphere at each stage of the RK scheme improved the overall accuracy of the method.

The trajectory reconstruction procedure is repeated for $j = 1, \ldots, N$ giving a set of N departure points at time $t = t_m$, which we denote by $\Xi^m = \{\boldsymbol{\xi}_j^m\}_{j=1}^N$ and refer to as the *Lagrangian point set*. It is at these points that we need to compute q. However, q is only (approximately) known at the nodes in X and in general $X \neq \Xi^m$, for any $m > 0$. Thus, we need to interpolate q from X to Ξ^m.

Remark 1. *SL methods are not subject to the advective CFL condition for the purpose of numerical stability (on tracer transport problems). So, in principle, if the velocity field is known for all time and space, then one can use integrators such as MATLAB's ode45 method to take one time-step over the whole simulation time. However, for problems that feature forcing terms or that require solving for the velocity field during the simulation [18, 20, 21, 52], one may be forced to take smaller intermediate time-steps. Additionally, for complex flows with significant variability in small scales, time-step restrictions arise from conditions on the inverse of the flow Jacobian in order to produce meaningful results [53]. This can result in time-step restrictions similar to the advective CFL condition. In problems where the flow fields are given and are smooth and slowly varying (as considered in this study), it is common to use a fixed time-step and (high-order) time integrator such that spatial errors dominate (see, for example, [54]). This is the approach we follow.*

3.3 Interpolation

Let $\mathbf{q}_X^m$ denote the vector containing the approximate solution to (1) at the Eulerian node set X (with the same ordering as the nodes in X) at time $t = t_m$. Below we discuss the details of interpolating $\mathbf{q}_X^m$ to the Lagrangian point set Ξ^m using the three different RBF interpolation methods presented in Section 2. These methods can all be used in the interpolation phase of Algorithm 1.

3.3.1 Global RBF interpolation

The most straightforward interpolation scheme is built on the global RBF interpolant (2). To compute this interpolant, we first solve the linear system (3), with the right hand side set equal to $\mathbf{q}_X^m$, to determine the interpolation coefficients $\mathbf{c}_X$. We then evaluate the interpolant at the points in Ξ^m, resulting in a vector of approximate values, which we denote by $\mathbf{q}_\Xi^m$. Finally, we set $\mathbf{q}_X^{m+1} = \mathbf{q}_\Xi^m$ to complete a time step.

The coefficient matrix A_X in (3) does not change throughout the simulation (since the nodes X are fixed for all time), thus allowing us to factorize A_X once as a preprocessing step. We use the inverse multiquadric (IMQ) radial kernel, $\phi(r) = (1 + (\varepsilon r)^2)^{-1/2}$, in this study, which is positive-definite, so that the matrix A_X can be factorized using Cholesky factorization. Since A_X is dense, this requires $O(N^3)$ operations initially. Each time step then requires $O(N^2)$ operations to solve (3) for each $\mathbf{q}_X^m$. The evaluation of (2) at Ξ^m requires an additional $O(N^2)$ computations. While this cost becomes prohibitive for large N, the global RBF method remains competitive for smooth solutions because of its high accuracy—faster than any algebraic order when the target function is sufficiently smooth [35]—as illustrated in the numerical results section.

3.3.2 Local RBF interpolation

The local RBF interpolation scheme is more complicated to implement than the global case. The first step of this scheme requires determining the n-point local stencil from X that will be used for interpolating to each of the points in Ξ^m. The approach we take is to determine the nearest neighbor in X for each $\boldsymbol{\xi}_j^m \in \Xi^m$, $j = 1, \ldots, N$. The nearest neighbor determines the center point for the stencil that will be used for the

interpolation. Letting $\mathcal{K}_j$, denote the index of the nearest neighbor in X for $\boldsymbol{\xi}_j$, the interpolation stencil is then denoted as $X_{\mathcal{K}_j}$. After these stencils have been determined, we perform the interpolations to the points in Ξ^m. These interpolations are done, for each $\boldsymbol{\xi}_j^m$, $j = 1, \ldots, N$, using (4), with $k = \mathcal{K}_j$, which results in the vector of approximate values $\mathbf{q}_\Xi^m$. As in the global case, we then set $\mathbf{q}_X^{m+1} = \mathbf{q}_\Xi^m$ to complete the time step.

The computation of the local interpolant for $\boldsymbol{\xi}_j$ requires solving the linear system (5) with the right hand side set equal to $\tilde{\mathbf{q}}_{X_{\mathcal{K}_j}}$. While it may be possible for some points in Ξ^m to share the same stencil, in general there will be $O(N)$ of these linear systems to solve. The matrices for these linear systems can all be pre-computed and factorized as a pre-processing step since the points in X do not change and all the stencils can be determined *prior* to the simulation. Unlike the global case, we use LU-factorizations of the matrices, since they are no longer positive-definite due to the inclusion of the polynomial terms and the use of the PHS kernel. The method initially requires $O(n^3 N)$ operations to factorize the linear systems for each stencil, then the cost is $O(n^2 N)$ operations to determine the interpolation coefficients of the local interpolants for all the points in Ξ^m and to evaluate the local interpolants to obtain $\mathbf{q}_\Xi^m$. To determine the nearest neighbors for each $\boldsymbol{\xi}_j^m \in \Xi^m$ in X, we use a kd-tree of the nodes X. This requires $O(N \log N)$ operations to construct initially, and $O(\log N)$ operations for each nearest neighbor search on average. Since n is chosen independent of N and typically $n << N$, the method has an asymptotic computational cost of $O(N \log N)$ per time-step, with a one-time initial cost of $O(N \log N)$.

Even when including the cost of kd-tree searches for nearest neigbhors, the local method is much more computationally efficient than the global RBF method. With this approach, it is possible to obtain high-order methods with algebraic convergence rates that depend on the stencil size n and the degree of the appended spherical harmonics.

3.3.3 RBF-PU interpolation

The first step of the RBF-PU interpolation scheme is to construct the local interpolants s_ℓ in (9) over the patches (spherical caps) Ω_ℓ, $\ell = 1, \ldots, M$. This requires solving M linear systems of the form (5), where the node sets and right hand sides used are determined by which nodes from X are included in patch Ω_ℓ. Once all the local interpolants are constructed, they are combined into the globally-smooth interpolant (9), which is then evaluated at all the points in Ξ^m. Note that for a given $\boldsymbol{\xi}_j^m \in \Xi^m$, the sum in (9) only needs to be taken over terms corresponding to the patches Ω_ℓ containing $\boldsymbol{\xi}_j$. As with the global and local methods, the result is a vector of approximate values $\mathbf{q}_\Xi^m$, which are assigned to $\mathbf{q}_X^{m+1}$ to complete the time step.

The computational complexity of the RBF-PU scheme is determined by how the patches Ω_ℓ, $\ell = 1, \ldots, M$, are distributed; this also directly effects the accuracy of the interpolants. Since the node sets X are assumed to be quasi-uniformly distributed, it makes sense to distribute the patches in a quasi-uniform manner to control the computational cost. Distributing the patches in a quasi-uniform manner is equivalent to distributing the set of patch centers $\{\boldsymbol{\omega}_\ell\}_{\ell=1}^M$ in a quasi-uniform manner. For these sets we use minimum energy (ME) points, which are computed by arranging the points in the set such that the Reisz energy (with a power of 2) of the set on $\mathbb{S}^2$ attains a minimum [46]. We use the pre-computed quasi-ME point sets from [49], which are available for $2 \leq M \leq 5000$.

What remains to determine the distribution of patches is M, the total number of patches (centers), and R_ℓ, $\ell = 1, \ldots, M$, the radii of the patches. We determine these quantities by setting 1) the number nodes, n, that each patch is to approximately contain and 2) the average number of patches, a, that a given node is to belong to. The first of these controls the computational cost for each patch and effects the accuracy of the local interpolants s_ℓ, as each interpolant will be based on approximately n nodes. The second value a controls the computational cost of evaluating the global RBF-PU interpolant (9), as it determines how much overlap there is amongst the patches. This value also controls the locality of the interpolant.

We use n to determine the radii of the patches as follows. If X contains N quasi-uniformly distributed nodes and there are to be approximately n nodes per patch, then the area per node over the entire sphere should approximately equal the area per node over patch Ω_ℓ, i.e. $4\pi/N \approx \pi R_\ell^2/n$, where R_ℓ is the radius of Ω_ℓ. This gives the estimate $R_\ell \approx 2\sqrt{n/N}$. Since the centers of the patches are also assumed to be quasi-uniformly distributed, the radii can all be chosen in the same way. We thus set one radius,

$$R = 2\sqrt{n/N}, \tag{11}$$

9

for every patch. We use a to determine the number of patches M as follows. Since the patch centers $\{\boldsymbol{\omega}_\ell\}_{\ell=1}^{M}$ are quasi-uniformly distributed, the same area arguments as above can be used to arrive at the estimate $4\pi/M \approx \pi R^2/a$. Solving for M in this equation and using (11), we obtain the following approximation (noting that M should be an integer)

$$M = \lceil aN/n \rceil. \tag{12}$$

Figure 2 illustrates the patches for $N = 4096$, $n = 100$, and $a = 2.5$.

Having defined these quantities, we can now estimate the computational cost of the RBF-PU interpolation scheme (per time-step). Similar to the local RBF scheme, the matrices in the M linear systems (5) for determining s_ℓ, $\ell = 1, \ldots, M$ do not change per time step and can thus be LU decomposed as a preprocessing step. Since each linear system contains approximately n nodes and there are M linear systems, this cost is $O(n^3 M)$, which can be estimated as $O(an^2 N)$ from (12). Using the LU factorizations, the M linear systems (5) can be solved in $O(anN)$ operations per time step. The evaluation of any of the interpolants s_ℓ in (9) at a point on $\mathbb{S}^2$ requires $O(n^2)$ operations. Each point of Ξ^m will, on average, belong to a patches, so that the sum in (9) will, on average, only involve a non-zero terms. Thus, the computational cost of evaluating the RBF-PU interpolant (9) at Ξ^m is $O(an^2 N)$. To determine which patches a given point in Ξ^m belongs to, we use a kd-tree of the patch centers, which only has to be constructed once at a cost of $O(M \log M) = O(aN/n \log(aN/n))$ operations and can then be searched each time-step at a cost of $O(N \log(aN/n))$. Since n and a are chosen independent of N and typically $a << n << N$, the method has an asymptotic computational cost of $O(N \log N)$ per time-step, with a one time initial cost of $O(N \log N)$.

4   Results

We investigate various properties of the three new methods proposed in this article on three standard test cases for transport on the sphere from the literature: solid-body rotation of a cosine bell from [26] and deformational flow with two different initial conditions from [27]. We test our methods using quasi-uniformly distributed node sets X of various sizes N. For the global RBF method, we use maximum determinant node sets [47] of sizes $N = 3136, 4096, 5041, 6084, 7744, 9025, 10000, 11881, 13689$, and 15129. For the local RBF and RBF-PU methods, we use equidistributed icosahedral nodes [55] of sizes $N = 2562, 5762, 10242, 23042, 40962$, and 92162. All these node sets are available from [49]. In the results that follow, we estimate the convergence of our method verses $\sqrt{N}$ since this is roughly inversely proportional to the average spacing between the nodes (see Section 3.1 for a discussion). For the local RBF and RBF-PU methods we present results for $n = 17, 31, 49$, and 84. These are common values used in RBF-FD methods for advective PDEs on the sphere [6, 7], and thus provide easy comparisons with those methods.

The main focus of the investigation is on accuracy of the methods, which we measure using the relative ℓ_2 and ℓ_∞ norms. We also give results on the dissipation and dispersion errors using a visual inspection of the errors and the *a posteriori* quantitative measures derived in [56] (see also [57]). These measures are derived from a decomposition of the mean square error of the numerical solutions and are given as follows:

$$\text{Dissipation error: } [\sigma(q) - \sigma(q_X)]^2 + [\bar{q} - \bar{q_X}]^2, \tag{13}$$

$$\text{Dispersion error: } 2\left[\sigma(q)\sigma(q_X) - \frac{1}{4\pi}\int_{\mathbb{S}^2}(q - \bar{q})(q_X - \bar{q_X})dS\right], \tag{14}$$

where q is the exact solution and q_X is the approximate solution, and bars on these variables denote the mean, while σ denotes their standard deviation over $\mathbb{S}^2$. In the results, we divide these errors by the mean square error of the solution to give a relative measure. Finally, we present results on the conservation properties of the methods by plotting the absolute errors in the total mass: $\left|\frac{1}{4\pi}\int_{\mathbb{S}^2}(q - q_X)dS\right|$. All of the above quantities are computed using discrete approximations to the continuous operators defining them on $\mathbb{S}^2$, which requires computing numerical approximations to surface integrals over $\mathbb{S}^2$ using the nodes X. To compute these surface integrals, we use the sixth-order kernel-based meshfree quadrature method from [58].

While not reported here, we also experimented with other quasi-uniform node sets for the global RBF method, specifically the the ME and spherical t-design nodes. In [2] (see also [59]), it was found that the MD vastly outperform the ME nodes when using flat RBFs (close to the spherical harmonic limit [36]). For

this study, we used shape parameters, ε, that are not in the flat regime and found the the resulting errors for all the test cases considered below were roughly comparable amongst the different types of node sets, with the ME nodes giving slightly smaller errors overall. We chose the MD nodes since there are many more of these node sets available for the sizes we were interested in testing. Finally, it should be noted that we did not try and optimize the value of ε for each test problem and each node set, so there is potential room for improvement in some of the results presented for the global RBF method. However, optimization of ε is rarely possible in real applications.

4.1   Solid body rotation of a cosine bell

As a first test problem, we consider the standard Test Case 1 from Williamson et al. [26]. The components of the steady velocity field for this test case in spherical coordinates ($-\pi \leq \lambda \leq \pi$, $-\pi/2 \leq \theta \leq \pi/2$) are given by

$$u(\lambda, \theta) = \sin(\theta) \sin(\lambda) \sin(\alpha) - \cos(\theta) \cos(\alpha), \quad v(\lambda, \theta) = \cos(\lambda) \sin(\alpha).$$

This velocity field results in solid body rotation at an angle of α with respect to the equator. In all of our tests, we use one of the standard choices of $\alpha = \pi/2$, which corresponds to advection of the initial condition directly over the north and south poles. Since our method works purely in Cartesian coordinates, we use a change of basis to obtain the same velocity field in these coordinates. The initial condition is taken as a compactly supported cosine bell centered at $(1, 0, 0)$:

$$q(\boldsymbol{x}, 0) = \begin{cases} \frac{1}{2} \left(1 + \cos \left(\frac{\pi r}{R_{\mathrm{b}}} \right) \right) & \text{if } r < R_{\mathrm{b}}, \\ 0 & \text{if } r \geq R_{\mathrm{b}}, \end{cases} \tag{15}$$

where $\boldsymbol{x} = (x, y, z)$, $r = \arccos(x)$, and the support is set as $R_{\mathrm{b}} = 1/3$. This initial condition has a jump in the second derivative at $r = R_{\mathrm{b}}$, which makes the test susceptible to both dispersive and diffusive errors. The test calls for simulating advection of the initial condition to the final time of $T = 2\pi$, which corresponds to one full revolution of the bell over the sphere. For the convergence tests, we set the time step to $\Delta t = \pi/10$, which is necessary so that spatial errors dominate for all the finest node sets. On the finest node set for the global RBF method, this time-step gives a CFL number of approximately 12, while for the finest node set for the local RBF and RBF-PU methods this gives a CFL number of approximately 28.

Convergence results in relative ℓ_2 and ℓ_∞ norms for increasing N are shown in Figure 3 for the three methods, together with the estimated rates of convergence. From the top row of this plot, we see that the ℓ_2 error for the local RBF and RBF-PU methods appears to converges at a rate close to 2.5 for all n, while the convergence for the global RBF method appears closer to 3. The ℓ_∞ error is shown in the bottom row and we see that the convergence rates are now around 2 for the local RBF and RBF-PU methods and 2.5 for the global RBF method. These rates of convergence are dictated by the smoothness of the initial condition, which is only $C^1(\mathbb{S}^2)$, and are consistent with the convergence results for the global RBF collocation and RBF-FD methods for this same problem [1, 6]. For the local RBF and RBF-PU methods, we see that increasing n leads to a decrease in the error, but not in the convergence rates (again because of the limited smoothness of the solution). We also see that the errors for the global method are smaller for similar values of N.

Figure 4 displays the dissipation and dispersion errors (13)–(14). Here we have fixed N at 40962 for the local RBF and RBF-PU methods and plotted dissipation and dispersion errors against with n, the local stencil/patch size. Also included in the plots are the results for the global RBF method (dashed line) with $N = 15129$. From the plots, we see that dispersion errors dominate the numerical solutions for all the methods, which is expected since the initial condition is only $C^1(\mathbb{S}^2)$. We also see that for the local RBF and RBF-PU methods increasing n (which increases the order of accuracy of these methods) leads to a decrease in both dissipation and dispersion errors, with the decrease in the former being much more pronounced.

We further explore the dissipation/dispersion errors of the methods by plotting the difference in the exact and approximate solutions after one ($T = 2\pi$) and ten ($T = 20\pi$) revolutions of the cosine bell. This test was suggested in [6] to examine whether the errors remain well-localized to the support of the bells over time, and to give a visual indication of the dissipation/dispersion errors. Figure 5 displays the results for the global RBF method for the case of N=15129 and the local RBF and RBF-PU methods, both for the

11

Figure 3: Convergence results of the relative ℓ_2 (top row) and ℓ_∞ (bottom row) errors for the solid body rotation of a cosine bell test case after one revolution. The results for the global RBF method are included in all plots for comparison purposes. The dashed lines are the lines of best fit to the data (without the first point included) of the form $C_n N^{-p_n/2}$. The values of p_n, which estimate the order of accuracy of the methods, are listed in the legend and the subscript on p is dropped for the global RBF method.

case of $N = 40962$ and $n = 49$. These values were selected since the errors were of similar magnitude. We see from the first row of the figure (one revolution) that the errors for the local RBF and RBF-PU methods are localized around the discontinuities in the second derivative of the solution, with the error being lower and more localized for the RBF-PU method. The errors for the global RBF method are spread more over the support of the entire bell. None of the methods display a dispersive wave-train over the whole sphere. After ten revolutions (second row of the figure) the errors for the local RBF and RBF-PU methods increase by about a factor of 4, but still remain localized around the discontinuities in the solution. The RBF-PU method again displays better localization of the errors, or lower dispersive-type errors. The dominant errors for the global method also remain restricted to the support of the bell, but only increase by about a factor of 1.5. The results for the local RBF and RBF-PU methods are qualitatively similar to those in [6] for the

12

Figure 4: Relative dissipation (left) and dispersion (right) errors (13)–(14) for the solid body rotation of a cosine bell test case after one revolution as n increases in the local RBF and RBF-PU methods. The global RBF method does not have a dependence on n and is included as a dashed line for reference.


(a) Error global RBF: $T = 2\pi$ (b) Error local RBF: $T = 2\pi$ (c) Error RBF-PU: $T = 2\pi$

(d) Error global RBF: $T = 20\pi$ (e) Error local RBF: $T = 20\pi$ (f) RBF-PU: $T = 20\pi$

Figure 5: Pseudocolor plots of the errors, $q_X - q$, on the sphere for solid body rotation test case of the cosine bell. The top row shows the errors after one revolution ($T = 2\pi$), while the bottom row shows the errors after ten revolutions ($T = 20\pi$). (a) & (d) show the results for the global RBF method using $N = 15129$, (b) & (d) show the results for the local RBF method using $N = 40962$ and $n = 49$, and (c) & (e) show the results for the RBF-PU method using $N = 40962$ and $n = 49$. Note the color scales differ for the two values of T. The view for all plots is from the initial center of the bell.

RBF-FD method with hyperviscosity stabilization.

Figure 6: Time traces of the mass conservation errors over ten revolutions of the cosine bell in the solid body rotation test case. For the local RBF and RBF-PU methods $N = 40962$ and $n = 49$, while for the global method $N = 15129$. For this test we used $\Delta t = \pi/20$.

Time traces of the absolute value of the mass conservation errors over ten revolutions ($T = 20\pi$) of the cosine bell are displayed in Figure 6. We again use $N = 40962$ and $n = 49$ for the local RBF and RBF-PU methods and $N = 15129$ for the global method. From this figure we see that the conservation errors for the local RBF and RBF-PU methods increase slightly over the integration time, with the growth of the local method being larger, while the errors in the global method remain consistent over the integration interval. While not presented here, we found that increasing n to 84 for the local RBF and RBF-PU methods had a marginal effect on decreasing the conservation errors.

4.2 Deformational flow

For the second test problem, we consider the deformational flow test case from [60]. The components of the velocity field for this test are given in spherical coordinates as

$$u(\lambda, \theta, t) = \frac{10}{T} \cos\left(\frac{\pi t}{T}\right) \sin^2\left(\lambda - \frac{2\pi t}{T}\right) \sin(2\theta) + \frac{2\pi}{T} \cos\theta, \tag{16}$$

$$v(\lambda, \theta, t) = \frac{10}{T} \cos\left(\frac{\pi t}{T}\right) \sin\left(2\lambda - \frac{2\pi t}{T}\right) \cos(\theta), \tag{17}$$

where $T = 5$. This flow field is designed to deform the initial condition up to time $t = 2.5$ and then reverse so that the solution is returned to its initial position and value at time $t = 5$. This value serves as the final time for the simulation. As before, we use a simple change of basis to convert this velocity field into Cartesian coordinates.

The following two initial conditions are considered:

1. Two cosine bells: $q(\boldsymbol{x}, 0) = 0.1 + 0.9(q_1(\boldsymbol{x}, 0) + q_2(\boldsymbol{x}, 0))$, where for $j = 1, 2$

$$q_j(\boldsymbol{x}, 0) = \begin{cases} \frac{1}{2}\left(1 + \cos\left(2\pi \cos^{-1}\left(\boldsymbol{x} \cdot \boldsymbol{p}_j\right)\right)\right) & \text{if } \cos^{-1}\left(\boldsymbol{x} \cdot \boldsymbol{p}_j\right) < \frac{1}{2}, \\ 0 & \text{otherwise.} \end{cases}$$

2. Two Gaussian bells: $q(\boldsymbol{x}, 0) = 0.95\left(e^{-5\|\boldsymbol{x}-\boldsymbol{p}_1\|_2^2} + e^{-5\|\boldsymbol{x}-\boldsymbol{p}_2\|_2^2}\right)$.

In both cases, $\boldsymbol{p}_1 = \left(\frac{\sqrt{3}}{2}, \frac{1}{2}, 0\right)$ and $\boldsymbol{p}_2 = \left(\frac{\sqrt{3}}{2}, -\frac{1}{2}, 0\right)$. Like the previous example, the first initial condition is a good test of the sensitivity of the three methods to dispersion errors as it is only $C^1(\mathbb{S}^2)$. The second

14

initial condition is $C^\infty(\mathbb{S}^2)$ and hence is a good test of the maximum convergence rates of the three methods. For all experiments involving the local RBF and RBF-PU methods, we set $\Delta t = 1/10$. This gives a CFL number of approximately 27 on the finest node set. For the global RBF method we use $\Delta t = 1/10$ for the cosine bells and $\Delta t = 1/40$ for the Gaussian bells, which results in CFL numbers of 12 and 3, respectively, on the finest node sets. All time-steps were chosen so that spatial errors dominate for all values of N.

4.2.1 Results for the cosine bells

Figure 7 displays the convergence results in relative ℓ_2 and ℓ_∞ norms for increasing N for the three methods, together with the estimated rates of convergence. We see that the convergence rates for this test case are a bit higher than the solid body rotation test, but that the errors are larger for a given N. The rates of convergence are again limited by the smoothness of the solution. The larger errors for this test are expected as the solution undergoes much more dramatic changes over the simulation period than the solid body rotation test. As in the solid body rotation test, the errors for the local RBF and RBF-PU methods are larger than the global method for a given N, but because the former methods are more computationally efficient, we can push them to larger N and reach smaller overall errors. At the end of this section, we compare the accuracy of all three methods to their runtime costs to get a better picture of their overall efficiency.

The relative dissipation and dispersion errors (13)–(14) are displayed in Figure 4. As in the previous test, N is fixed for the local RBF and RBF-PU methods and the errors are plotted against n. The global RBF method with $N = 15129$ is displayed as a dashed line. The figure shows that dispersion errors again dominate the numerical solutions for all the methods and that increasing n for the local RBF and RBF-PU methods leads to a decrease in these errors. Visual depictions of these dispersive errors are given in Figure 9. Part (a) shows the initial condition and final solution for this test case and parts (b)–(d) show the difference between this solution and the numerical solutions at the final time for the three methods. The error for the global method (part (b)) are about 1.5 times higher than the local RBF and RBF-PU methods and appears to be much more dispersive in nature. Also, the error for the RBF-PU method is more localized to the discontinuities of the initial condition than both the global and local RBF methods.

Time traces of the mass conservation errors for the three methods over the simulation time are displayed in Figure 10 (a). For the local RBF and RBF-PU methods, we see that, after a relatively large initial growth, the errors level off around $t = 1$ and then start to grow very slowly towards the end of the simulation. The global RBF error is overall larger than the other two methods, but does not exhibit a discernible growth rate.

4.2.2 Results for the Gaussian bells

The convergence results for the relative ℓ_2 and ℓ_∞ errors for this test are displayed in Figure 11. We see from these plots that the norm of the errors for global RBF method appear to converge faster than any polynomial rate, which is expected since the solution is $C^\infty(\mathbb{S}^2)$ (see Section 2.1). The convergence rates of the errors for the local RBF and RBF-PU methods are also higher for this smooth test case. Unlike the two previous tests, we see that these convergence rates also increase as n increases. The RBF-PU method appears to have a higher convergence rate for the same n than the local RBF method, and the errors for the RBF-PU method are lower for each corresponding n and N value. This is likely due to the global smoothness of the RBF-PU interpolant.

Figure 12 displays the relative dissipation and dispersion errors just like the other test cases. We see that the dissipation and dispersion errors are smallest for the global method and largest for the local RBF method. We also see that increasing n for local and RBF-PU methods leads to a much larger decrease in both the dissipation and dispersion errors than the previous two test cases. We omit plots of the errors over the sphere since the dispersion errors are not a real issue for this test case.

Time traces of the mass conservation errors are displayed in Figure 10 (b). As with the dissipation/dispersion errors, the global method shows the best results followed by the RBF-PU method and then the local RBF method. All of these methods show a slight growth in the conservation errors towards the end of the simulation time. While not reported here, increasing n in the local and RBF-PU methods does lead to a decrease in the conservation errors.
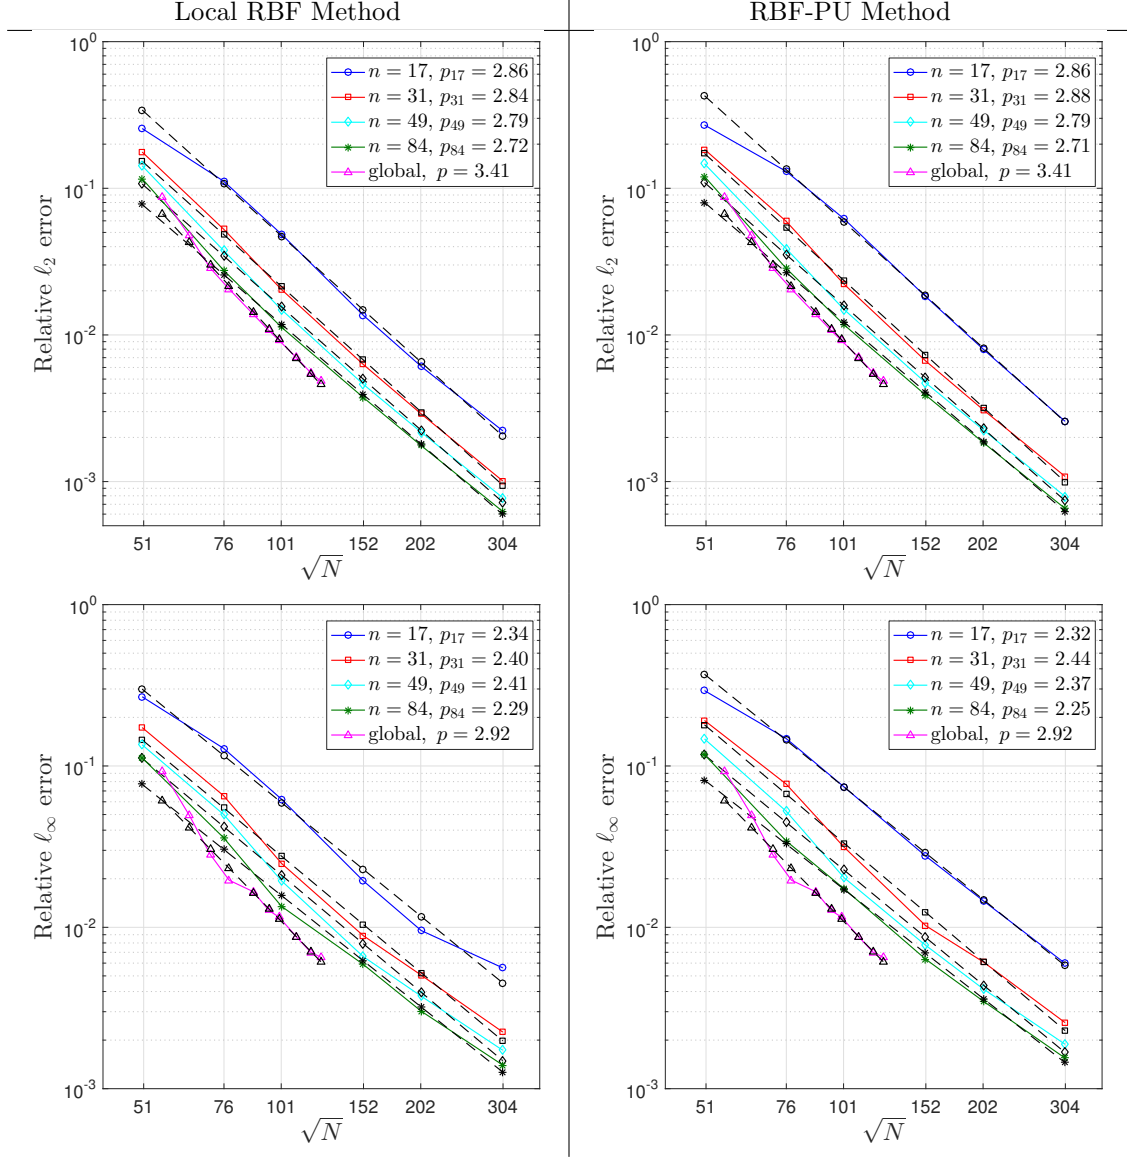
Figure 7: Convergence results of the relative ℓ_2 (top row) and ℓ_∞ (bottom row) errors the deformational flow of two cosine bells test case. The results for the global RBF method are included in all plots for comparison purposes. The dashed lines are the lines of best fit to the data (without the first point included) of the form $C_n N^{-p_n/2}$). The values of p_n, which estimate the order of accuracy of the methods, are listed in the legend and the subscript on p is dropped for the global RBF method.

4.3 Cost versus accuracy

To properly compare all three methods, it is important to look at their cost (measured in wall-clock time) versus accuracy. In this section we give such a comparison for the deformational flow test case using both the non-smooth and smooth initial conditions. We examine the simulation time (ignoring preprocessing costs) and accuracy of each method (in the relative ℓ_2 norm), and how these change with N and n. All tests were run on a Linux workstation with a 3 GHz Intel Core i7-3930K (12 logical cores) and 32 GB of RAM. All the codes were written and executed in MATLAB (version 2016a) in standard double precision, but without any explicit parallelization. Unlike the previous sections, we here vary the time-steps with the number of nodes N to approximately optimize the overall efficiency of the methods. The time-steps used for the various node

Figure 8: Relative dissipation (left) and dispersion (right) errors (13)–(14) for the deformational flow test case of two cosine bells after one revolution as n increases in the local RBF and RBF-PU methods. The global RBF method does not have a dependence on n and is included as a dashed line for reference.

Local RBF and RBF-PU						
N	2562	5762	10242	23042	40962	92162
CB, Δt	5/20	5/25	5/30	5/35	5/40	5/45
GB, Δt	5/20	5/40	5/60	5/80	5/100	5/120

Global RBF									
N	3136	4096	5041	6084	7744	9025	10000	11881	13689	15129
CB, Δt	5/20	5/20	5/25	5/25	5/30	5/35	5/35	5/40	5/40	5/45
GB, Δt	5/20	5/40	5/60	5/80	5/100	5/120	5/140	5/160	5/180	5/200

Table 1: Time-steps used for the numerical results in Sections 4.3 and 4.4 for the deformational flow test cases of the cosine bells (CB) and Gaussian bells (GB).

sizes are listed in Table 1. As mentioned previously, these time-steps were selected so that spatial errors in all the methods dominate.

The results for the non-smooth cosine bells are displayed in Figure 13 (a). We see that for low accuracy, the global RBF method has a lower overall cost than both the local RBF method and for the RBF-PU method with $n = 17$ and $n = 31$. For higher accuracies, the cost of the global method increases much more rapidly than the other two methods because of its $O(N^2)$ complexity and, for high enough accuracies, both the local RBF and RBF-PU methods are more efficient overall for all values of n. Comparing the local RBF and RBF-PU methods, we see that, for a fixed n, the latter has a lower computational cost for a given accuracy. Also, for this non-smooth test case, it does not appear to offer much benefit in terms of cost to use large n with the local method, whereas the RBF-PU method shows slightly better efficiency with increasing n.

The results for the smooth Gaussian bells test case are given in Figure 13 (b). These plots clearly illustrate the advantage of the global RBF method with a smooth kernel over the local RBF method and the RBF-PU method with smaller n when used on a problem with a smooth solution. Extrapolating out, we see that global method is able to reach a much smaller ℓ_2 error for same runtime for both the local and RBF-PU methods. However, the global method will exhaust the memory resources of the machine much more quickly than the local and RBF-PU methods, and its results depend more heavily on memory latency. The figure also shows that for a smooth initial condition, it pays to use larger values of n, with this being much

(a) Exact solution: T=0 & 5　　(b) Error global RBF, $T = 5$

(c) Error local RBF: $T = 5$　　(d) Error RBF-PU: $T = 5$

Figure 9: (a) Pseudocolor plots on the sphere of the initial and final solutions for the two cosine bells test. (b)–(c) Show pseudocolor plots of the error, $q_X - q$, at $T = 5$ for the global RBF, local RBF, and RBF-PU methods, respectively. For the global RBF method, $N = 15129$, while for the local and RBF-PU methods, $N = 40962$ and $n = 49$. Note that the color scale is about 1.5 times larger for the global RBF method. The view for all plots is from the midpoint between the centers of the two bells.

(a) Cosine bells　　(b) Gaussian bells
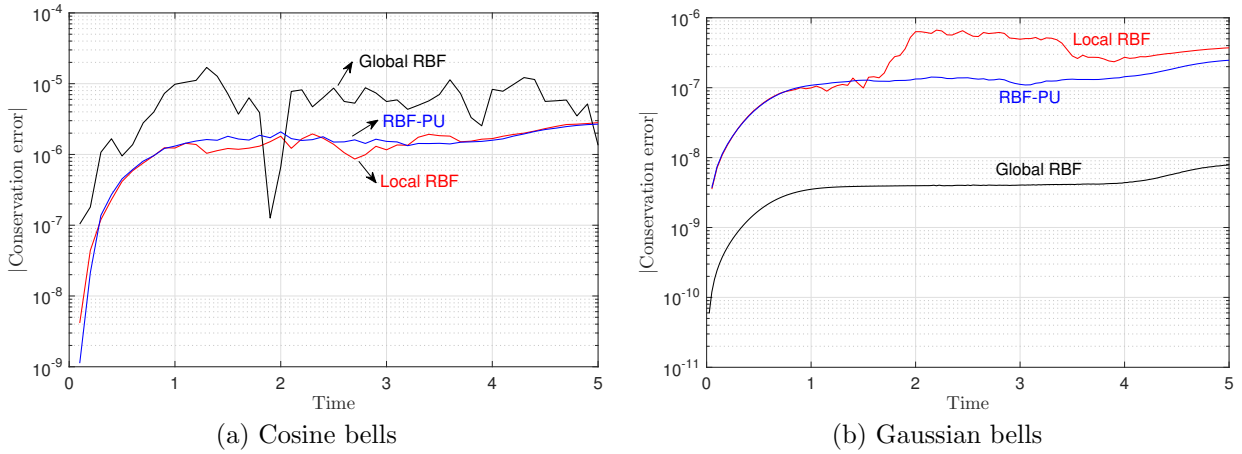
Figure 10: Time traces of the mass conservation errors over the simulation time for the deformational flow tests. For the local RBF and RBF-PU methods $N = 92162$ and $n = 49$ and $\Delta t = 1/10$ for both part (a) and (b), while for the global method $N = 15129$ and $\Delta t = 1/10$ for part (a) and $N = 15129$ and $\Delta t = 1/40$ for part (b).

more beneficial for the RBF-PU method. Comparing the local and RBF-PU methods for this smooth initial condition, we see that overall, in our implementation, the RBF-PU method again gives higher accuracy for
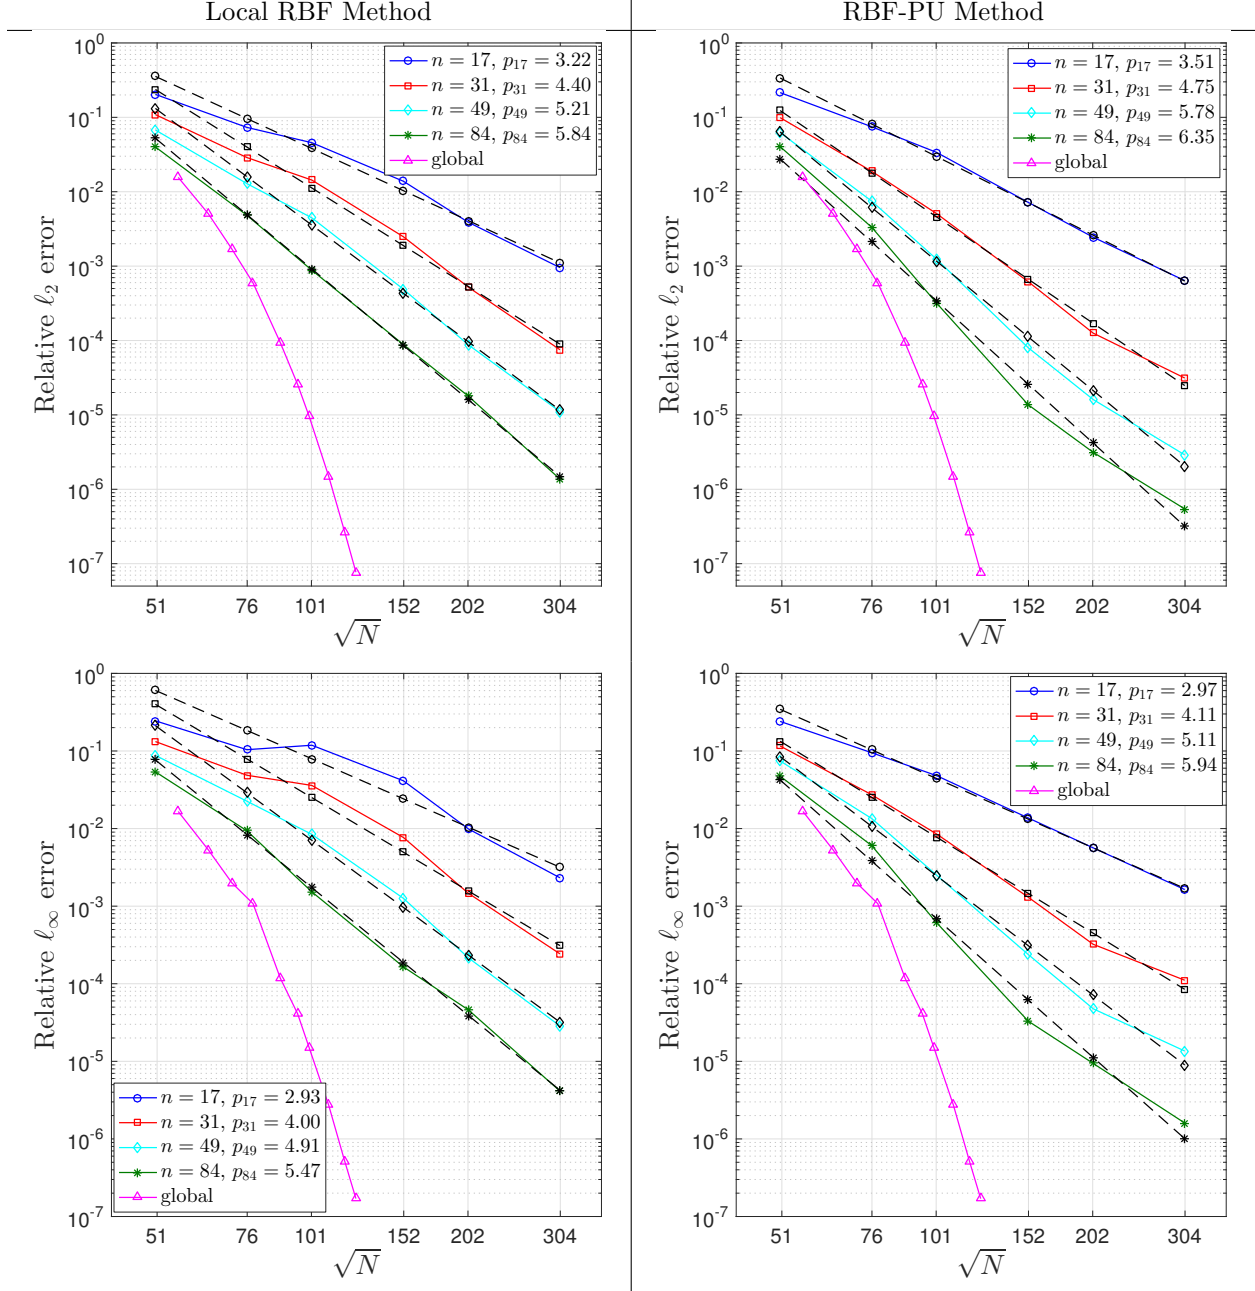
Figure 11: Convergence results of the relative ℓ_2 (top row) and ℓ_∞ (bottom row) errors the deformational flow of two Gaussian bells test case. The results for the global RBF method are included in all plots for comparison purposes. The dashed lines are the lines of best fit to the data (without the first point included) of the form $C_n N^{-p_n/2}$. The values of p_n, which estimate the order of accuracy of the methods, are listed in the legend. Lines of best fit are omitted for the global case since the convergence appears to be faster than any polynomial order.

the same cost.

It is important to note that these are serial implementations of all the methods. The local RBF method is easily parallelized on SIMD architectures, allowing for speedups of 2-8 times over a serial implementation, as shown for the related RBF-FD method in [8, 9]. The RBF-PU method also has promising parallelization properties using its patch-based structure. In contrast, the global RBF method with infinitely-smooth kernels

Figure 12: Relative dissipation (left) and dispersion (right) errors (13)–(14) for the deformational flow test case of two Gaussian bells after one revolution as n increases in the local RBF and RBF-PU methods. The global RBF method does not have a dependence on n and is included as a dashed line for reference.

requires a domain-decomposition style approach to parallelization. We thus expect parallel versions of both the local and RBF-PU methods to give much better results in cost versus accuracy studies than presented here. We also note that, in all our tests, the global and RBF-PU methods were able to utilize MATLAB's automatically multithreaded BLAS kernels more efficiently than the local RBF method, since these two methods inherently use fewer, denser matrix-matrix multiplications (BLAS-3) operations.

4.4 Comparison with other methods

In this section, we compare the three new SL RBF methods against other methods for transport on the sphere found in the literature. First, we compare the overall efficiency (wall-clock time vs. accuracy) of the SL RBF methods against the standard Eulerian RBF-FD method developed in [6] (see also [7, 8, 9]) on the deformational flow of two cosine bells test case. For the results presented here, we used the Gaussian RBF, similar to [6, 7], with an $n = 84$ point stencil. For the shape parameter ε, we used the approach suggested in [7] and selected ε to scale with N as $\varepsilon = 0.063\sqrt{N}$. As mentioned in the introduction, this method requires a hyperviscosity stabilization term of the form $(-1)^{k+1}\gamma_N\Delta^k$, where Δ is the Laplacian, k is a positive integer, and γ_N is a constant that decreases with N. With the choices of RBF, ε, and n above, we found that selecting the hyperviscosity parameters as $k = 6$ and $\gamma_N = 15 * 2^{-8}N^{-5}$ produced stable results with the time-steps used, without overdamping the solutions. As in [6, 7], we used RK4 as the time-integrator and node sets of size $N = 5762, 10242, 23042, 40962$, and 92162, with corresponding time-steps $\Delta t = 5/400, 5/700, 5/900, 5/1400$, and $5/2800$. These are approximately the largest time-steps that could be used so that the solutions were stable and spatial errors dominate. For the SL RBF methods, we used the time-steps listed in Table 1. The same machine described in the previous section was used for all the simulations.

Figure 14 shows the results of the comparison between the SL RBF methods and the Eulerian RBF-FD method. We see that the local RBF and RBF-PU methods have a lower error for the same degrees of freedom and that they require less computational time than the Eulerian RBF-FD method. This increase in efficiency comes from the fact that much larger time-steps can be used in the SL simulations.

For the second comparison, we also use the deformational flow test, but now for both the cosine bells and the Gaussian bells. The focus here is on the errors for similar resolutions, degrees of freedom, and time-steps. We also include three commonly used methods from the geosciences (CSLAM, DG, and HOMME) in the comparison. Table 2 displays the results and the caption contains descriptions of the methods used. We can

20

Method	Resolution (in degrees)	Degrees of freedom (N)	Time-step (Δt), non-dimensional	Relative ℓ_2 error
Cosine Bells				
CSLAM [61]	$0.75°$	86400	5/240	$\approx 6 \cdot 10^{-3}$
DG, $p = 3$ [27]	$1.5°$	38400	5/2400	$1.39 \cdot 10^{-2}$
RBF-FD, $n = 84$ [6]	$1.5°$	23042	5/900	$1.17 \cdot 10^{-2}$
Local RBF, $n = 84$	$1.5°$	23042	5/35	$3.45 \cdot 10^{-3}$
RBF-PU, $n = 84$	$1.5°$	23042	5/35	$3.63 \cdot 10^{-3}$
Global RBF	$1.64°$	15129	5/45	$5.1 \cdot 10^{-3}$
Gaussian Bells				
CSLAM [61]	$0.75°$	86400	5/240	$\approx 5 \cdot 10^{-4}$
HOMME, $p = 6$ [54]	$1.5°$	29400	5/4800	$\approx 3 \cdot 10^{-3}$
RBF-FD, $n = 84$ [6]	$1.5°$	23042	5/900	$3.18 \cdot 10^{-4}$
Local RBF, $n = 84$	$1.5°$	23042	5/80	$5.50 \cdot 10^{-5}$
RBF-PU, $n = 84$	$1.5°$	23042	5/80	$1.35 \cdot 10^{-5}$
Global RBF	$1.64°$	15129	5/200	$7.68 \cdot 10^{-8}$

Table 2: Comparison of various methods from the literature for the two deformational flow test cases. CSLAM is the conservative semi-Lagrangian multi-tracer transport scheme used in [61], which is based on a cubed sphere grid. DG is the discontinuous Galerkin scheme used in [27]. This is an Eulerian scheme that uses $p = 3$ degree polynomials (fourth-order accurate) and the cubed sphere grid. HOMME is the High-Order Methods Modeling Environment scheme used in [54]. This is also an Eulerian scheme using a cubed-sphere grid and the results given are for a continuous Galerkin formulation using $p = 6$ degree polynomials. RBF-FD is the mesh-free Eulerian scheme from [6] using $n = 84$ point stencils and the Gaussian RBF with the shape-parameter selected using the formulas from [7]. Values appearing with a $\approx$ were obtained from plots of the relative ℓ_2 errors given in the corresponding referenced papers, as no exact values are reported. For the RBF-FD method the results were generated by the present authors using the code from [7]. The results given for the CSLAM, DG, and HOMME methods are for the non-filtered (or non shape-preserving) versions, which give the lowest errors for these test cases. Resolution (in degrees) is an approximate measure of the spacing of the nodes (or grid points) around the equator.

(a) Deformational flow: Cosine Bells

(b) Deformational flow: Gaussian Bells

Figure 13: Computational cost (runtime) verses accuracy for the deformational flow test case using (a) Cosine Bells and (b) Gaussian bells. The left column shows the results for the local RBF method, while the right is for the RBF-PU method. The global RBF results are repeated in all plots for comparison purposes. The runtime is for the total simulation and does not include any pre-computational costs.

see from the table that all three SL RBF methods compare quite favorably for these test cases to existing methods. It should be noted, however, that the CSLAM, DG, and HOMME methods are all mass-conserving, whereas, the present version of the RBF SL methods presented here are not.
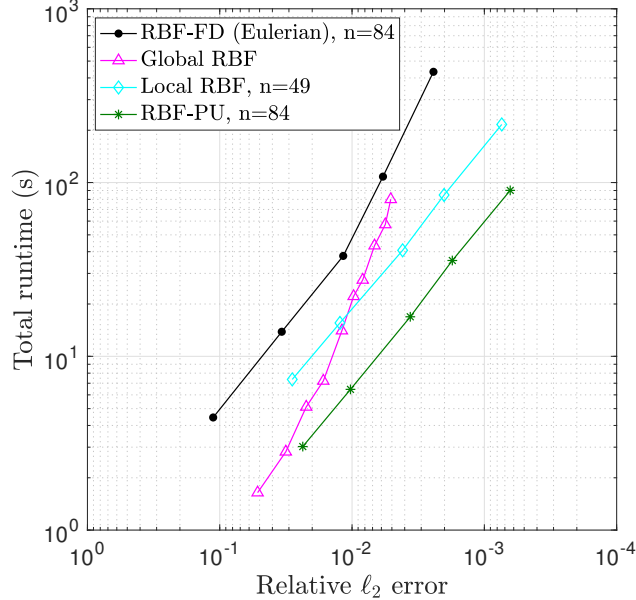
Figure 14: Comparison of the three new SL RBF schemes and the Eulerian RBF-FD scheme from [6]. The plot shows the computational cost (runtime) verses relative ℓ_2 error for the Cosine Bells deformational flow test case. The stencil sizes for the global RBF, local RBF, RBF-PU, and Eulerian RBF-FD methods are shown in the legend. For the latter three methods, the same node sets were used of sizes $N = 5762, 10242, 23042, 40962$, and 92162.

5   Summary

In this article, we presented three new SL methods for transport on a sphere based on interpolation with global RBFs, local RBFs, and RBF-PU. The RBF framework allowed us to obtain either high-order convergence rates for smooth problems using only scattered nodes and Cartesian coordinates. Using scale-free RBFs with the addition of spherical harmonics in the local and RBF-PU methods removes the need to choose a shape parameter and also avoids stagnation errors observed in other applications of these methods. Additionally, the SL framework appears to lend our methods intrinsic stability without the need for hand-tuned artificial hyperviscosity. We summarize the features of our methods below:

- The global RBF method with smooth kernels appears to have the best cost-accuracy profile for problems with smooth solutions that can be resolved using a relatively small number of nodes. However, the gap between the global method and local/RBF-PU methods can be decreased by increasing the stencil/patch sizes, n. Furthermore, we expect that the local/RBF-PU methods will be more competitive for problems with smooth solutions with fine structures that need to be resolved with large numbers of points.

- For non-smooth solutions, the local and RBF-PU methods both out perform the global method for moderate to high accuracies.

- The RBF-PU method gives smaller errors, lower dispersion, and better conservation properties than the local RBF method for a comparable number of degrees of freedom, regardless of the smoothness of the solutions.

- The local RBF method is easier to implement and may be more amenable to highly efficient implementations on SIMD architectures.

- All three methods compare favorably to commonly used methods for transport on the sphere for the problems considered here.

Our algorithms were designed for transport in an incompressible velocity field. However, compressible fields arise naturally in many biological and geophysical problems. We are currently working on extending our method to handle this case. In addition, there is a need for quasi-monotone (non-oscillatory) local RBF methods to ameliorate the numerical dispersion seen therein. We plan to address this issue in future work. Furthermore, the local RBF method can obtain very large speedups even in a serial implementation if used with the overlapped RBF-FD framework [42]; we plan to explore this in future work as well. Finally, the RBF framework allows for a straightforward generalization of our methods to more general manifolds than spheres [62, 63, 64]. While the methods discussed here are limited to transport on the sphere, we anticipate that they may be extended to the full non-linear shallow water equations using a similar approach as [20].

Acknowledgements

References

[1] N. Flyer, G. B. Wright, Transport schemes on a sphere using radial basis functions, J. Comput. Phys. 226 (2007) 1059–1084.

[2] B. Fornberg, C. Piret, On choosing a radial basis function and a shape parameter when solving a convective pde on a sphere, J. Comput. Phys. 227 (2008) 2758–2780.

[3] N. Flyer, G. B. Wright, A radial basis function method for the shallow water equations on a sphere, Proc. Roy. Soc. A 465 (2009) 1949–1976.

[4] Q. T. L. Gia, Approximation of parabolic PDEs on spheres using spherical basis functions, Adv. Comput. Math. 22 (2005) 377–397.

[5] G. B. Wright, N. Flyer, D. A. Yuen, A hybrid radial basis function-pseudospectral method for thermal convection in a 3-D spherical shell, Geochem. Geophys. Geosyst. 11 (2010) Q07003.

[6] B. Fornberg, E. Lehto, Stabilization of RBF-generated finite difference methods for convective PDEs, J. Comput. Phys. 230 (2011) 2270–2285.

[7] N. Flyer, E. Lehto, S. Blaise, G. B. Wright, A. St-Cyr, A guide to RBF-generated finite differences for nonlinear transport: shallow water simulations on a sphere, J. Comput. Phys. 231 (2012) 4078–4095.

[8] E. Bollig, N. Flyer, G. Erlebacher, Solution to PDEs using radial basis function finite-differences (RBF-FD) on multiple GPUs, J. Comput. Phys. 231 (2012) 7133–7151.

[9] M. Tillenius, E. Larsson, E. Lehto, N. Flyer, A scalable RBF-FD method for atmospheric flow, J. Comput. Phys. 298 (2015) 406–422.

[10] N. Flyer, G. B. Wright, B. Fornberg, Radial Basis Function-generated Finite Differences: A Mesh-free Method for Computational Geosciences, in: W. Freeden, M. Z. Nashed, T. Sonar (Eds.), Handbook of Geomathematics, 2nd Edition, Springer-Verlag, Berlin, 2014, pp. 1–30.

[11] K. A. Aiton, A Radial Basis Function Partition of Unity Method for Transport on the Sphere, Master's thesis, Boise State University, USA (2014).

[12] A. Safdari-Vaighani, A. Heryudono, E. Larsson, A radial basis function partition of unity collocation method for convection–diffusion equations arising in financial applications, J. Sci. Comput. 64 (2015) 341–367.

[13] N. Flyer, G. A. Barnett, L. J. Wicker, Enhancing finite differences with radial basis functions: Experiments on the Navier-Stokes equations, J. Comput. Phys. 316 (2016) 39–62.

[14] A. Staniforth, J. Côté, Semi-Lagrangian integration schemes for atmospheric models—a review, Mon. Wea. Rev. 119 (1991) 2206–2223.

[15] A. Staniforth, N. Wood, Aspects of the dynamical core of a nonhydrostatic, deep-atmosphere, unified weather and climate-prediction model, J. Comput. Phys. 227 (2008) 3445–3464.

[16] P. K. Smolarkiewicz, J. A. Pudykiewicz, A class of semi-Lagrangian approximations for fluids, J. Atmos. Sci. 49 (1992) 2082–2096.

[17] P. K. Smolarkiewicz, L. G. Margolin, On forward-in-time differencing for fluids: an Eulerian/semi-Lagrangian non-hydrostatic model for stratified flows, Atmos.-Ocean 35 (1997) 127–152.

[18] D. Xiu, G. E. Karniadakis, A semi-Lagrangian high-order method for Navier–Stokes equations, J. Comput. Phys. 172 (2001) 658–684.

[19] P. J. Rasch, D. L. Williamson, On shape-preserving interpolation and semi-Lagrangian transport, SIAM J. Sci. and Stat. Comput. 11 (1990) 656–687.

[20] A. T. Layton, W. F. Spotz, A semi-Lagrangian double fourier method for the shallow water equations on the sphere, J. Comput. Phys. 189 (2003) 180–196.

[21] M. A. Tolstykh, Vorticity-divergence semi-Lagrangian shallow-water model of the sphere based on compact finite differences, J. Comput. Phys. 179 (2002) 180–200.

[22] R. Nair, J. Côté, A. Staniforth, Cascade interpolation for semi-Lagrangian advection over the sphere, Q. J. Roy. Meteor. Soc. 125 (1999) 1445–1468.

[23] P. H. Lauritzen, R. D. Nair, P. A. Ullrich, A conservative semi-Lagrangian multi-tracer transport scheme (CSLAM) on the cubed-sphere grid, J. Comput. Phys. 229 (2010) 1401–1424.

[24] M. F. Carfora, Semi-Lagrangian advection on a spherical geodesic grid, Int. J. Numer. Method. Fluid. 55 (2007) 127–142.

[25] N. Flyer, B. Fornberg, V. Bayona, G. A. Barnett, On the role of polynomials in RBF-FD approximations: I. Interpolation and accuracy, J. Comput. Phys. 321 (2016) 21–38.

[26] D. L. Williamson, J. B. Drake, J. J. Hack, R. Jakob, P. N. Swarztrauber, A standard test set for numerical approximations to the shallow water equations in spherical geometry, J. Comput. Phys. 102 (1992) 211–224.

[27] R. D. Nair, P. H. Lauritzen, A class of deformational flow test cases for linear transport problems on the sphere, J. Comput. Phys. 229 (2010) 8868–8887.

[28] D. Y. Le Roux, C. A. Lin, A. Staniforth, An accurate interpolating scheme for semi-Lagrangian advection on an unstructured mesh for ocean modelling, Tellus A 49 (1997) 119–138.

[29] J. Behrens, A. Iske, Grid-free adaptive semi-Lagrangian advection using radial basis functions, Comput. Math. Appl. 43 (2002) 319–327.

[30] M. Zerroukat, A simple mass conserving semi-Lagrangian scheme for transport problems, J. Comput. Phys. 229 (2010) 9011–9019.

[31] B. Fornberg, N. Flyer, A Primer on Radial Basis Functions with Applications to the Geosciences, SIAM, Philadelphia, 2014.

[32] G. E. Fasshauer, Meshfree Approximation Methods with MATLAB, Interdisciplinary Mathematical Sciences - Vol. 6, World Scientific Publishers, Singapore, 2007.

[33] S. Hubbert, T. M. Morton, L_p-error estimates for radial basis function interpolation on the sphere, J. Approx. Theory 129 (2004) 58–77.

[34] F. J. Narcowich, X. Sun, J. D. Ward, H. Wendland, Direct and inverse Sobolev error estimates for scattered data interpolation via spherical basis functions, Found. Comput. Math. 7 (2007) 369–390.

[35] K. Jetter, J. Stöckler, J. D. Ward, Error estimates for scattered data interpolation on spheres, Math. Comput. 68 (1999) 733–747.

[36] B. Fornberg, C. Piret, A stable algorithm for flat radial basis functions on a sphere, SIAM J. Sci. Comput. 30 (2007) 60–80.

[37] Q. T. L. Gia, I. H. Sloan, H. Wendland, Multiscale analysis in Sobolev spaces on the sphere, SIAM J. Numer. Anal. 48 (2010) 2065–2090.

[38] G. B. Wright, B. Fornberg, Scattered node compact finite difference-type formulas generated from radial basis functions, J. Comput. Phys. 212 (2006) 99–123.

[39] H. Wendland, Fast evaluation of radial basis functions: Methods based on partition of unity, in: Approximation Theory X: Wavelets, Splines, and Applications, Vanderbilt University Press, 2002, pp. 473–483.

[40] R. Cavoretto, A. DeRossi, Fast and accurate interpolation of large scattered data sets on the sphere, J. Comput. Appl. Math. 234 (2010) 1505–1521.

[41] R. Cavoretto, A. De Rossi, E. Perracchione, Partition of unity interpolation on multivariate convex domains, Int. J. Model. Simul. Sci. Comp. 06 (2015) 1550034.

[42] V. Shankar, The overlapped radial basis function-finite difference (RBF-FD) method: A generalization of RBF-FD, J. Comput. Phys. 342 (2017) 211–228.

[43] V. Bayona, N. Flyer, B. Fornberg, G. A. Barnett, On the role of polynomials in RBF-FD approximations: II. Numerical solution of elliptic PDEs, J. Comput. Phys. 332 (2017) 257–273.

[44] J. A. Reeger, B. Fornberg, Numerical quadrature over the surface of a sphere, Stud. Appl. Math. 137 (2016) 174–188.

[45] M. Falcone, R. Ferretti, Semi-Lagrangian Approximation Schemes for Linear and Hamilton–Jacobi Equations, SIAM, Philadelphia, PA, 2013.

[46] D. P. Hardin, E. B. Saff, Discretizing manifolds via minimum energy points, Notices Amer. Math. Soc. 51 (2004) 1186–1194.

[47] R. S. Womersley, I. H. Sloan, How good can polynomial interpolation on the sphere be?, Adv. Comput. Math. 23 (2001) 195–226.

[48] J. R. Baumgardner, P. O. Frederickson, Icosahedral discretization of the two-sphere, SIAM J. Numer. Analy. 22 (1985) 1107–1115.

[49] G. B. Wright, SpherePts, https://github.com/gradywright/spherepts/ (2018).

[50] P. A. Bosler, J. Kent, R. Krasny, C. Jablonowski, A Lagrangian particle method with remeshing for tracer transport on the sphere, J. Comput. Phys. 340 (2017) 639–654.

[51] E. Fehlberg, Classical fifth-, sixth-, seventh-, and eighth-order Runge-Kutta formulas with stepsize control, Tech. Rep. TR R-287, NASA (1968).

[52] W. Bao, S. Jin, P. A. Markowich, Numerical study of time-splitting spectral discretizations of nonlinear Schrödinger equations in the semiclassical regimes, SIAM J. Sci. Comput. 25 (2003) 27–64.

[53] J.-F. Cossette, P. K. Smolarkiewicz, P. Charbonneau, The monge-ampère trajectory correction for semi-Lagrangian schemes, J. Comput. Phys. 274 (2014) 208–229.

[54] P. H. Lauritzen, P. A. Ullrich, C. Jablonowski, P. A. Bosler, D. Calhoun, A. J. Conley, T. Enomoto, L. Dong, S. Dubey, O. Guba, A. B. Hansen, E. Kaas, J. Kent, J.-F. Lamarque, M. J. Prather, D. Reinert, V. V. Shashkin, W. C. Skamarock, B. Sørensen, M. A. Taylor, M. A. Tolstykh, A standard test case suite for two-dimensional linear transport on the sphere: results from a collection of state-of-the-art schemes, Geosci. Model Dev. 7 (2014) 105–145.

[55] T. Michaels, Equidistributed icosahedral configurations on the sphere, Comput. Math. Appl. 74 (2017) 605–612.

[56] L. L. Takacs, A two-step scheme for the advection equation with minimized dissipation and dispersion errors, Mon. Wea. Rev. 113 (1985) 1050–1065.

[57] M. Restelli, L. Bonaventura, R. Sacco, A semi-Lagrangian discontinuous Galerkin method for scalar advection by incompressible flows, J. Comput. Phys. 216 (2006) 195–215.

[58] E. Fuselier, T. Hangelbroek, F. Narcowich, J. Ward, G. Wright, Kernel based quadrature on spheres and other homogeneous spaces, Numer. Math. 127 (2014) 57–92.

[59] B. Fornberg, J. M. Martel, On spherical harmonics based numerical quadrature over the surface of a sphere, Adv. Comput. Math. 40 (2014) 1169–1184.

[60] R. D. Nair, P. H. Lauritzen, A class of deformational flow test cases for linear transport problems on the sphere, J. Comput. Phys. 229 (2010) 8868–8887.

[61] P. H. Lauritzen, W. C. Skamarock, M. J. Prather, M. A. Taylor, A standard test case suite for two-dimensional linear transport on the sphere, Geosci. Model Dev. 5 (2012) 887–901.

[62] E. Fuselier, G. B. Wright, Order-preserving derivative approximation with periodic radial basis functions, Adv. Comput. Math. 41 (2015) 23–53.

[63] V. Shankar, G. B. Wright, R. M. Kirby, A. L. Fogelson, A radial basis function (RBF)-finite difference (FD) method for diffusion and reaction–diffusion equations on surfaces, J. Sci. Comput. 63 (2014) 745–768.

[64] E. Lehto, V. Shankar, G. B. Wright, A radial basis function (RBF) compact finite difference (FD) scheme for reaction-diffusion equations on surfaces, SIAM J. Sci. Comput. 39 (2017) A2129–A2151.