On choosing a radial basis function and a shape parameter when solving a convective PDE on a sphere

Bengt Fornberg [*,1], Cécile Piret [2]

University of Colorado, Department of Applied Mathematics, 526 UCB, Boulder, CO 80309, USA

Abstract

Radial basis function (RBF) approximations have been used for some time to interpolate data on a sphere (as well as on many other types of domains). Their ability to solve, to spectral accuracy, convection-type PDEs over a sphere has been demonstrated only very recently. In such applications, there are two main choices that have to be made: (i) which type of radial function to use, and (ii) what value to choose for their shape parameter (denoted by ε, and with flat basis functions – stretched out in the radial direction – corresponding to $\varepsilon = 0$). The recent RBF-QR algorithm has made it practical to compute stably also for small values of ε. Results from solving a convective-type PDE on a sphere are compared here for many choices of radial functions over the complete range of ε-values (from very large down to the limit of $\varepsilon \to 0$). The results are analyzed with a methodology that has similarities to the customary Fourier analysis in equispaced 1-D periodic settings. In particular, we find that high accuracy can be maintained also over very long time integrations. We furthermore gain insights into why RBFs sometimes offer higher accuracy than spherical harmonics (since the latter arise as an often non-optimal special case of the former). Anticipated future application areas for RBF-based methods in spherical geometries include weather and climate modeling.
© 2007 Elsevier Inc. All rights reserved.

1. Introduction

Many application areas, such as geophysics (including weather and climate modeling), astrophysics, quantum mechanics, etc. require PDEs to be solved in spherical geometries. Flyer and Wright [7,8] have recently shown that an RBF-based pseudospectral (PS) approach can be very successful for solving wave-type PDEs

---

[*] Corresponding author. Tel.: +1 303 492 5915; fax: +1 303 492 4066.
E-mail addresses: fornberg@colorado.edu (B. Fornberg), piret@colorado.edu (C. Piret).

(such as pure convection equations and the shallow water equations) on a spherical surface. Compared to other spectral approaches, such as double Fourier methods [13,28,29,35], spherical harmonics methods [2,19,21,37,39], and spectral element methods [18,30,36,38,40], they found the RBF approach to be particularly promising since it features

- algebraic simplicity (complete codes for PDE test cases are often less than 50 lines of Matlab),
- immediate generalizability from spherical to arbitrarily shaped surfaces, and
- opportunities for combining spectral accuracy with local refinement.

Most RBFs depend on a shape parameter ε, with $\varepsilon \to 0$ corresponding to the limit of increasing flatness. Lowering the value of ε usually increases the resulting accuracy to some point when one or both of the following two factors halt or reverse this trend:

(1) Onset of disastrous numerical ill-conditioning if using a direct implementation of the RBF procedure (denoted RBF-Direct, to be explained in Section 2.1).
(2) Onset of a Runge phenomenon reminiscent of that for polynomial interpolation, as explained in [17], Section 3 (potentially disastrous mainly in the presence of boundaries or variable node densities, else mainly setting a limit for the accuracy that can be reached).

Fig. 1 illustrates these two factors in three different cases from the literature. The top row of subplots shows how the error varies with ε when solving a Poisson equation over a unit disk, as described in [25]. The dashed comparison line shows the accuracy that is reached with the previously most accurate available procedure, Fourier pseudospectral (PS) in the angular direction and Chebyshev PS radially. In all cases, the resolution was 16 nodes on the boundary and 48 in the interior. In the top right subplot, computational ill-conditioning for small ε was eliminated with the Contour–Padé algorithm [15]. The middle row of subplots shows similar trends in the context of interpolating a Gaussian bell over the surface of a sphere when using $n = 1849$ nearly uniformly distributed nodes. For further details in this case, see [14]. The RBF-QR algorithm, introduced in that work, was used (instead of the Contour–Padé algorithm) to eliminate the ill-conditioning for the right subplot. The single subplot in the bottom row (with data and further explanations given in [7]) shows the max norm error after a cosine bell has been convected once around the sphere when solving the time dependent PDE (5), to be described later in Section 3. The main goal of the present study is to explore this last test case further, in particular when the numerical ill-conditioning for low ε-values has been eliminated. By means of both numerical computations and through some novel analysis, we will

(1) Explore how the accuracy varies with
  - RBF type,
  - Shape parameter (for full range of $\varepsilon \geqslant 0$),
  - Length of time integration.
(2) Discuss the $\varepsilon = 0$ case since this, as was found in [14], is usually equivalent to using spherical harmonics (SPH) in place of RBFs.

We present in Section 2 first a very brief introduction to RBF interpolation, and we then quote some relevant results from the literature, such as the significance of the flat basis function limit. The convective PDE and results using one type of RBF discretization are presented in Section 3. We extend in Section 4 these results to also include a large number of different RBF choices, again over the full range of ε-values. In the present context, the differences in performance between the smooth RBF types are found to be minor, whereas the performances of the non-smooth RBF are much inferior. Some of these results can be understood through the analysis in Section 5, generalizing to scattered node situations on the sphere the Fourier-based arguments that are routinely used on equispaced periodic problems in 1-D. In Section 6, focusing on summaries and conclusions, we again address why (in the present context of long time integration), smooth basis functions are superior to non-smooth ones, regardless of whether the solution that is convected is smooth or not. The Appendices A and B provide technical details on some RBF-QR implementation issues.
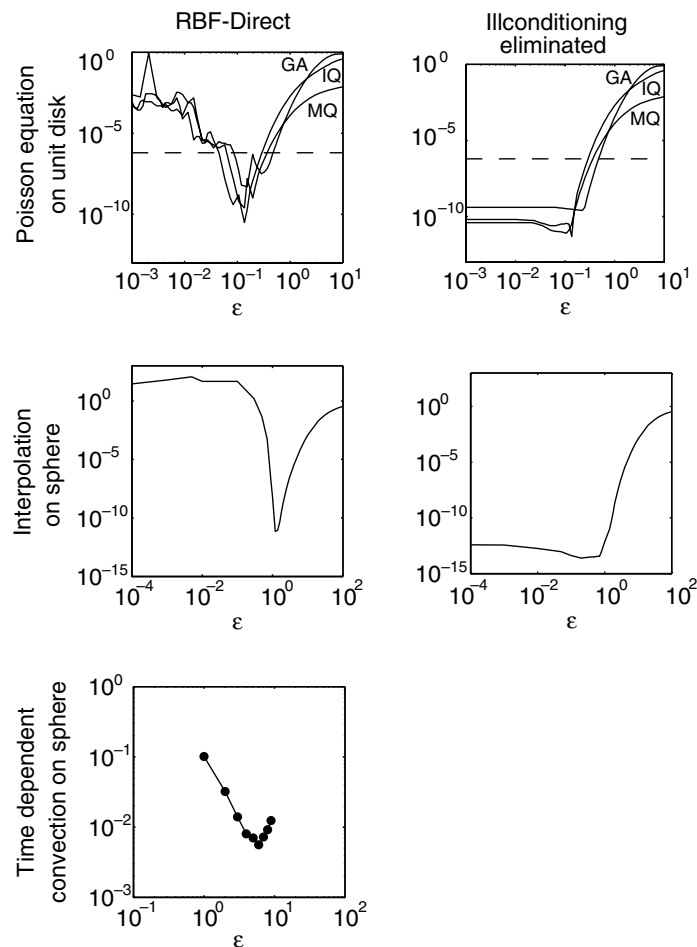
Fig. 1. RBF errors in three applications, displayed as functions of the shape parameter ε. The three rows of subplots reproduce data from [25,14,7], respectively. The abbreviations MQ, IQ, and GA are explained in Section 2. The present work will provide the missing bottom right subplot (cf. Fig. 17).

2. RBF methodology

2.1. The form of an RBF interpolant

The basic RBF interpolant takes the form

$$s(\underline{x}) = \sum_{i=1}^{n} \lambda_i \phi(\|\underline{x} - \underline{x}_i\|), \tag{1}$$

where $\|\cdot\|$ denotes the Euclidean norm. In order for it to take the values f_i at locations $\underline{x}_i$, $i = 1, 2, \ldots n$, the expansion coefficients λ_i need to satisfy

$$A\underline{\lambda} = \underline{f}, \tag{2}$$

where the entries of the matrix A are $A_{i,j} = \phi(\|\underline{x}_i - \underline{x}_j\|)$. We denote numerical use of (2) followed by (1) as "RBF-Direct". In this study, we will concentrate our attention on the radial functions $\phi(r)$ listed in Table 1. The parameter ε, included in all but the piecewise smooth global cases CU and TPS, is known as the *shape parameter*. The listed Wendland functions [42] are those of lowest degree which guarantee that the A-matrix is positive definite for all distinct node locations in 2 and 3 dimensions. This property holds in all dimensions also for IMQ, IQ, and GA. In the MQ, CU, and TPS cases, A is still symmetric. For MQ, non-singularity remains guaranteed, but positive definiteness is lost. Additional issues arise in the CU and TPS cases. Commonly used variations of (1) are for MQ

Table 1
Definitions of some types of radial functions

Name of RBF	Abbreviation	Definition
Smooth, global		
Multiquadric	MQ	$\sqrt{1+(\varepsilon r)^2}$
Inverse multiquadric	IMQ	$\frac{1}{\sqrt{1+(\varepsilon r)^2}}$
Inverse quadratic	IQ	$\frac{1}{1+(\varepsilon r)^2}$
Gaussian	GA	$e^{-(\varepsilon r)^2}$
Piecewise smooth, global		
Cubic	CU	$\lvert r \rvert^3$
Thin plate spline	TPS	$r^2 \ln \lvert r \rvert$
Piecewise smooth, compact (for $0 \leqslant r \leqslant \frac{1}{\varepsilon}$; equal to zero for $r > \frac{1}{\varepsilon}$)		
Wendland type, order 2	W2	$(1-\varepsilon r)^4(4\varepsilon r+1)$
Order 4	W4	$(1-\varepsilon r)^6\left(\frac{35}{3}(\varepsilon r)^2+6\varepsilon r+1\right)$
Order 6	W6	$(1-\varepsilon r)^8(32(\varepsilon r)^3+25(\varepsilon r)^2+8\varepsilon r+1)$

The shape parameter ε controls their 'flatness'. In the case of the Wendland functions, their 'order' refers to their degree of smoothness (C^2, C^4, C^6, respectively).

$$s(\underline{x}) = \alpha + \sum_{i=1}^{n} \lambda_i \phi(\lVert \underline{x} - \underline{x}_i \rVert), \ \text{with the constraint} \ \sum_{i=1}^{n} \lambda_i = 0 \tag{3}$$

and for CU and TPS (in the case of 3-D)

$$s(\underline{x}) = \alpha + \beta x + \gamma y + \delta z + \sum_{i=1}^{n} \lambda_i \phi(\lVert \underline{x} - \underline{x}_i \rVert) \tag{4}$$

with the constraints

$$\sum_{i=1}^{n} \lambda_i = \sum_{i=1}^{n} \lambda_i x_i = \sum_{i=1}^{n} \lambda_i y_i = \sum_{i=1}^{n} \lambda_i z_i = 0.$$

For more detailed discussions, see [3,32,34].

The most notable feature of the Wendland functions is their compact support in case ε is large, leading to sparse A-matrices, and therefore with possibilities for high computational speeds both in evaluating (1) and in solving (2) (for example if using conjugate gradient-type methods). In the present case with nodes on the unit sphere, all sparsity is lost if $\varepsilon < \frac{1}{2}$.

2.2. RBFs for PDEs

In 1990, Kansa introduced collocation with RBFs as a means to approximate spatial derivatives, and thus to numerically solve PDEs [22,23]. In case of smooth RBFs, this approach is typically spectrally accurate [4,27,46]. Another notable advantage with this RBF approach (compared to, say, finite difference, finite element, and finite volume methods) is that it replaces the often challenging task of creating computational meshes over irregular domains with the easier one of scattering computational nodes.

Driscoll and Fornberg [5] observed that, in the flat basis function limit $\varepsilon \to 0$ for globally smooth RBFs in 1-D, the interpolant in general converges to Lagrange's interpolation polynomial. From this follows that the RBF approach for PDEs can be viewed as a generalization (to irregular domains and to scattered nodes) of the pseudospectral (PS) method [1,10,41]. This approach is therefore nowadays often described as the RBF-PS method. In a 1-D periodic setting, the $\varepsilon \to 0$ limit will reproduce the Fourier-PS method [5]. For nodes on a sphere, the corresponding limit was found in [14] to agree with a spherical harmonics (SPH-PS) method. More results about RBF in the flat basis function limit can be found for example in [11,12,16,26,33].

As noted in the introduction, Flyer and Wright [7] were the first to use RBFs to solve purely convective (i.e. non-dissipative) PDEs over a spherical surface. Because they implemented their scheme based on direct use of (2) followed by (1), ill-conditioning prevented their numerical explorations from being extended also to arbitrarily small values of ε.

Much more materials on both 'RBF methodology' and 'RBFs for PDEs' can be found in the two recent books [6,42].

3. Time dependent PDE on a sphere

3.1. Test problem

The standard PDE test problem that we will consider describes 'solid body' rotation/convection around an axis that is inclined by the angle α relative to the polar axis (cf. [7,13,43], and Fig. 2a). Following the convention in many applications, we define spherical coordinates as shown in Fig. 2b

$$\begin{cases} x = \rho \cos \varphi \cos \theta \\ y = \rho \sin \varphi \cos \theta \; . \\ z = \rho \sin \theta \end{cases}$$

This differs from standard spherical coordinates (as used for ex. in [14]) both in the use of latitude (as opposed to co-latitude) and also in the directions denoted by θ and φ. In the present coordinates, the governing PDE becomes

$$\frac{\partial u}{\partial t} + (\cos \alpha - \tan \theta \sin \varphi \sin \alpha) \frac{\partial u}{\partial \varphi} - \cos \varphi \sin \alpha \frac{\partial u}{\partial \theta} = 0. \tag{5}$$

One full revolution will correspond to the time $t = 2\pi$.

The pole singularities enter through the $\tan \theta$ factor when $\theta = \pm\frac{\pi}{2}$. These singularities are not physical, but arise only as a consequence of the (θ, φ)-system itself being singular at these locations. When the RBF differentiation matrix (DM, representing the spatial operator of (5) in terms of the node values of u) is formed as described in [7], these singularities will therefore vanish (the DM reflects the physical operator, but not what coordinate system it happened to be expressed in during an intermediate derivation step).

The initial condition that we will use is the *Cosine bell*, shown in Fig. 3. It can be described in Cartesian coordinates as


Fig. 2. (a) Flow directions in the 'solid body' convection test problem. (b) Spherical latitude–longitude-type coordinate system.
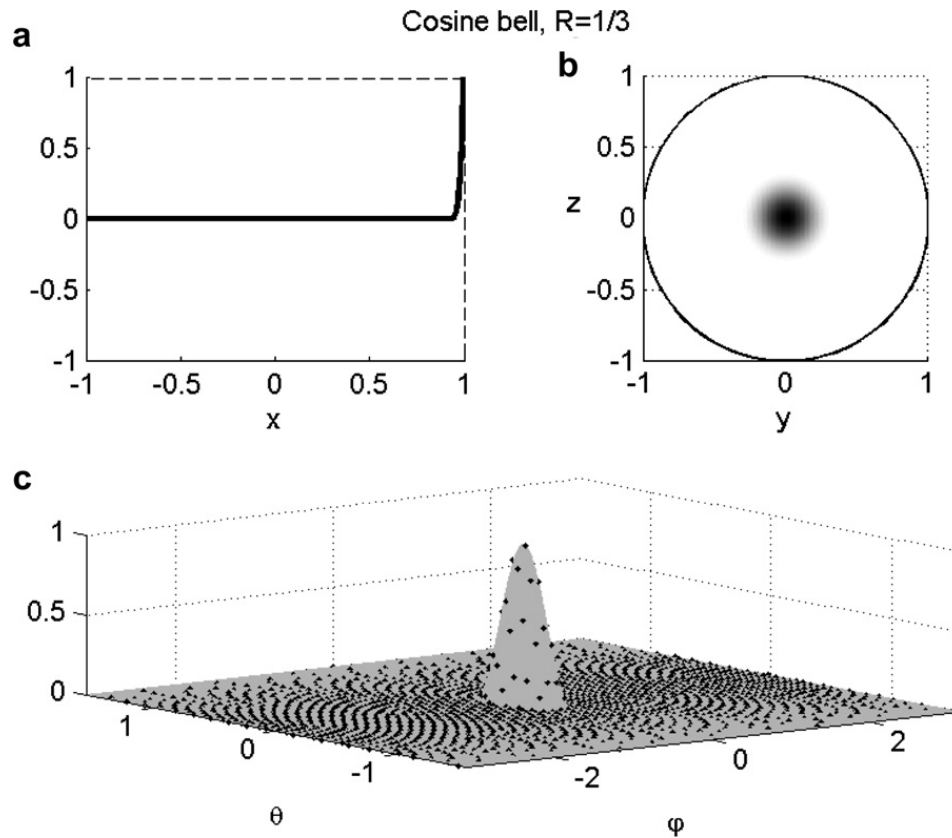
Fig. 3. Three illustrations of the cosine bell (a) as a function of x according to (6), (b) in grey-scale on sphere surface, viewed from positive x-direction, and (c) on an 'unrolled' φ, θ-plane (with the $n = 1849$ ME node locations also marked).

$$u(x,y,z,0) = \begin{cases} \frac{1}{2}\left(1 - \cos\left(\frac{\pi}{R}\arccos x\right)\right) & x > \cos R \\ 0 & x \leqslant \cos R \end{cases} \tag{6}$$

with $R = 1/3$, and restricted to the surface of the unit sphere.

Cosine bells have become standard initial conditions for convective test calculations on a sphere [21,30,35,38] for several reasons:

(1) Their compact support makes it particularly easy to display and to interpret dispersive errors after convection.
(2) The easy-to-change peak width R allows testing to be carried out on different spatial scales.
(3) The discontinuous second derivative at the base of the bell prevents very high order methods from producing misleadingly good results compared to what might be expected in cases of more physically relevant data.

3.2. Different node distributions

Because of the clear advantages seen in previous works [7,14] of using near-uniform rather than fully random node distributions, we limit the discussion here to the former case. When interpolating using SPH, Womersley and Sloan [44,45] noticed very large differences also between different types of near-uniform distributions. The two near-uniform node sets that are used in this study are both taken from [44], and are shown in Fig. 4. They are described as 'minimal energy' (ME) and 'maximal determinant' (MD), respectively. The former can be obtained as the equilibrium of freely moving nodes on the sphere surface, repelling each other like equal point charges. In the MD case, node locations are instead obtained by maximizing the determinant

B. Fornberg, C. Piret / Journal of Computational Physics 227 (2008) 2758–2780

ME nodes   MD nodes

Fig. 4. Illustration of the minimal energy (ME) and maximal determinant (MD) $n = 1849$ node sets.

of a certain 'Gram matrix'. Womersley and Sloan report that SPH interpolation errors (measured in any standard norm, or as a Lebesgue constant) can be several orders of magnitude larger in ME than in MD cases. Since RBF interpolation in the $\varepsilon \to 0$ limit usually corresponds to SPH interpolation, it is not surprising that the choice of nodes (ME vs. MD) will influence the RBF accuracy also for non-zero ε (although gradually less so when ε is increased, as will be illustrated in Section 3.4).

3.3. Method of lines (MOL) formulation and time stepping considerations for the PDE test problem

The test problem (5) can be written more briefly $\frac{\partial u}{\partial t} + L(u) = 0$, where L is the spatial derivative operator. The RBF-Direct-based DM that arises when approximating the spatial derivative terms in (5) was shown in [7] to take the form

$$D = B \cdot A^{-1}, \tag{7}$$

where B is antisymmetric, with elements

$$B_{i,j} = L\phi(||\underline{x} - \underline{x}_j||)|_{\underline{x}=\underline{x}_i} \tag{8}$$

and A is defined in (2), with elements

$$A_{i,j} = \phi(||\underline{x}_i - \underline{x}_j||). \tag{9}$$

We will here compute this matrix D either immediately through (7)–(9) (RBF-Direct), or via RBF-QR, as described in Appendix A. The ill-conditioning problem with RBF-Direct originates from the fact that, as $\varepsilon \to 0$, A becomes very nearly singular, causing elements of A^{-1} to diverge to plus or minus infinity. A vast amount of numerical cancellation then occurs when the O(1)-sized matrix D is formed as the matrix product (7).

The key concept behind the RBF-QR algorithm is the recognition that the basis functions $\phi(||\underline{x} - \underline{x}_i||)$, $i = 1, 2, \ldots, n$, when ε is small, form a terrible base that nevertheless spans an excellent approximation space. The RBF-QR algorithm expands each basis function in a SPH series, after which it transpires that we can obtain a very well-conditioned base *spanning exactly the same space*. Using this new base bypasses the ill-conditioning but leads otherwise to identical results. Explained somewhat differently: For all values of ε, the entries of D behave in a well-conditioned way with respect to perturbations in the input data, i.e. the node locations $\underline{x}_i$. In the case of small ε, RBF-Direct is an ill-conditioned algorithm for numerically computing D, while RBF-QR is a well-conditioned one.

A matrix that is a product of an antisymmetric and a positive definite one can only have purely imaginary eigenvalues. The DMs that arise from RBF discretization of the spatial operator in (5) will thus be of this type for the positive definite RBF choices in Table 1 (IMQ, IQ, GA, W2, W4, W6). Standard time stepping methods, such as the fourth order Runge–Kutta method (RK4), will therefore be applicable, no matter how the

nodes are scattered over the surface of the sphere. An extension to the result above, given in [31], extends it to (3) or (4) rather than (1) in the remaining cases of MQ, CU and TPS. In the present tests, we found the eigenvalues to be purely imaginary also without invoking this extension, allowing us to use (1) in all cases.

Numerical tests using RBFs for convecting the cosine bell (as well as results for the second test problem in [43]) were presented in [7]. The errors they reported after a computational time corresponding to one full revolution around the sphere ($t = 12$ days in their notation, $t = 2\pi$ in our notation) come from three different sources:

(1) Time stepping errors in the RK4 scheme.
(2) Properties of the DM.
(3) Errors when the numerical solution, defined at the grid points, is interpolated and compared against the analytic solution across the complete sphere.

The issue (1) can be resolved by using small enough RK4 time steps, so that the error types (2) and (3) will dominate (this still allows much longer time steps to be taken than what is feasible in the alternative methods, as reported in [7]; see also the discussion here in Section 5). In order to more clearly analyze the errors coming from the DM (issue 2), we will here make use of the fact that the discretized test problem can be integrated analytically in time. Letting $\underline{u}(t)$ denote a column vector with n entries, containing the numerical approximations at the n node points as functions of time, the discretization of (5) takes the form

$$\frac{\partial \underline{u}}{\partial t} + D\underline{u} = 0 \tag{10}$$

for which the analytic solution is

$$\underline{u}(t) = \mathrm{e}^{-tD}\underline{u}(0).$$

3.4. Numerical tests using IMQ RBF

The RBF-QR method is based on certain series expansions which involve powers of ε, and these might fail to converge for large ε. In contrast, RBF-Direct fails due to ill-conditioning for low values of ε. If the number of nodes n is relatively low (such as $n = 1849$ used throughout this study), there will usually be some range in which both methods work, but this overlap can get lost as n is increased further.

The first issue that arises (already at time $t = 0$, i.e. before time stepping has even started) is how accurately the cosine bell will be represented as an RBF interpolant of the discrete initial data that is shown in Fig. 3c. Fig. 5 (for SPH, i.e. for most smooth RBFs in their $\varepsilon \rightarrow 0$ limit) and the curve marked 'Interpolation' in Fig. 6a (for a wide range of ε values) show that direct interpolation based on the values at the $n = 1849$ ME nodes can lead to very large errors in-between the node points (positioned as seen in Figs. 3c and 4). This is entirely in agreement with the observations in [45]. The error level drops significantly when ε is increased from $\varepsilon = 0$, and it is particularly low around $2 \lesssim \varepsilon \lesssim 10$ (Fig. 6a). Only here is the error comparable to what is obtained with MD nodes throughout the full range $0 \leqslant \varepsilon \lesssim 10$ (Fig. 6b), confirming that the latter node type offers much more 'robust' interpolation. However, a still better option with regard to obtaining a uniformly good RBF representation of the initial condition is possible through a least squares approach, and is illustrated by the curves marked 'Least squares' in Fig. 6a–d. These RBF representations were obtained by a 2-stage process of first finding the best representation of the analytic initial condition (the cosine bell) in terms of SPH up to order $\mu = \sqrt{n} = 43$ (also featuring $\mu^2 = n = 1849$ free parameters) through least squares over a much denser point set, followed by evaluation of this SPH approximation back at the original $n = 1849$ node points. With this initial filter-type step, we obtain the initial condition across the full sphere with higher accuracy than was the case with immediate interpolation using either of the ME or MD node sets. Furthermore, the accuracy becomes almost completely independent of the details of the node distributions. Fig. 6c and d illustrates that the piecewise smooth RBF are much less sensitive in this regard (of ME vs. MD). In all the computations that are described below, we use ME nodes and this least square procedure to get the RBF

Fig. 5. The SPH interpolant to the cosine bell over the $n = 1849$ ME node points. The input data is as exactly as shown by the data points in Fig. 3c.
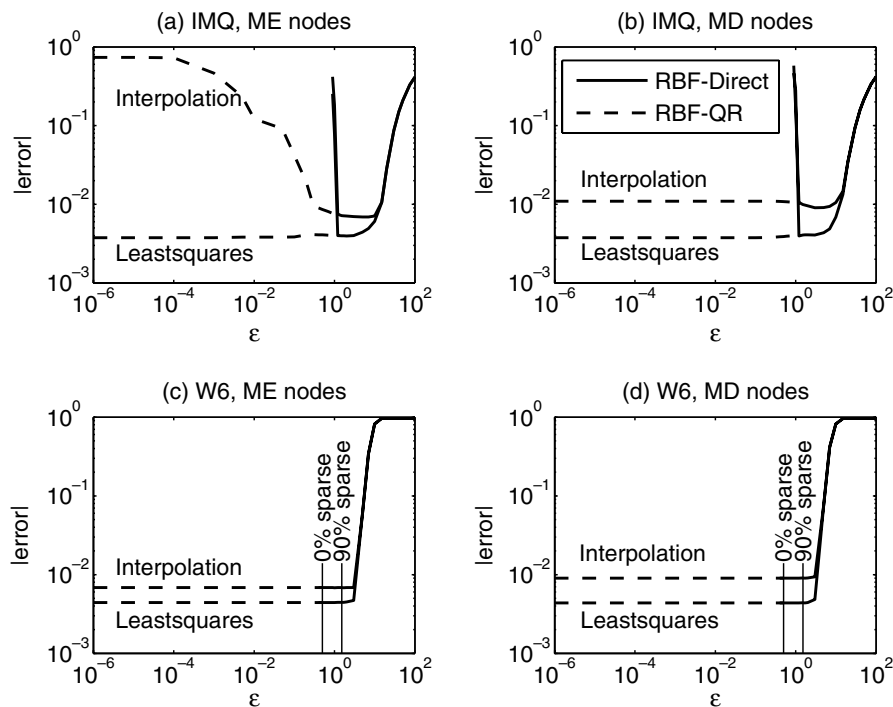
Fig. 6. The errors (in max norm) when the initial cosine bell is brought to an $n = 1849$ node RBF representation (a) IMQ on ME nodes, (b) IMQ on MD nodes, (c) W6 RBF on ME nodes, and (d) W6 on MD nodes. In the W6 cases, the ε values corresponding to 0% and 90% sparsity of the A-matrix are marked. In all cases, we contrast direct interpolation with the least square approach described in Section 3.4. We display the divergence of RBF-Direct for decreasing ε in subplots a and b (appearing as near-vertical solid lines around $\varepsilon = 1$). In subplots c and d and in the following Figs. 7, 9, and 10, similar divergence occurs but is not explicitly marked.

representation of the initial condition at $t = 0$. The same filter-type strategy as described above is a routine part of most SPH-PS calculations, then typically applied at every time step rather than as here only at $t = 0$.

Fig. 7a and b shows how the accuracy at $t = 10$ and $t = 10,000$ vary with ε in the cases of IMQ and W6 – typical results for a smooth and a piecewise smooth radial function, respectively. Including two more RBF types, Fig. 8a and b shows the time evolution of the error over $0 \leqslant t \leqslant 10$ and over $0 \leqslant t \leqslant 10,000$, respectively. We notice that the error degrades severely in time for non-smooth RBF, whereas it holds up extremely well for the smooth ones. In these and the following cases, the error was measured as the maximal discrepancy compared to the analytic solution, over the $n = 1849$ node points.

We extend next these comparisons to include all the radial functions listed in Table 1.

Fig. 7. Errors at times $t = 10$ and $t = 10,000$ as functions of ε in case of (a) IMQ and (b) W6. The thin dashed line shows the (ε-independent) error of the initial SPH representation of the initial data at $t = 0$.

Fig. 8. Evolution of the error with time over (a) $0 \leqslant t \leqslant 10$ and (b) $0 \leqslant t \leqslant 10,000$. The ε values for IMQ, GA and W6 were near zero, whereas TPS does not include any ε.

4. Comparisons between different RBF types

There has been numerous suggestions in the literature that some types of RBFs generally tend to give more accurate results than others. For example, MQ is often particularly highly regarded, whereas concerns have been raised against GA, sometimes for the reason that GA does not support any 'exact polynomial reproduction' properties on infinite lattices (when $\varepsilon > 0$) [3].

The accuracy that is achieved by different RBF methods is highly problem dependent. For example, interpolation of non-smooth data places very different demands on the method than does long-term solution of convective-type PDEs. With the combination of RBF-Direct and RBF-QR, we have now the ability to run the convective test problem for the full set of RBFs shown in Table 1, throughout the full range of ε (from zero and upwards), giving the results shown in Figs. 9 and 10. As we have noted already, smooth RBFs give excellent accuracy (not much different from the initial error at $t = 0$) also over very long time integrations. The differences between these smooth RBFs are in the present context seen to be very minor in comparison to the much more rapid loss of accuracy seen for the piecewise smooth RBF. These latter ones differ between themselves largely according to their smoothness, with the higher order ones (such as W4 and W6) being more accurate than, say, TPS, CU, or W2. In order to offer some accuracy over longer time integrations, also the Wendland functions need to be scaled so that they become relatively flat (ε small), causing their A-matrices to lose their sparseness. However, since the Wendland-based D-matrices (7) are never sparse, these RBFs do not (in the present PDE context) seem to offer any speed advantages over the standard smooth RBF types.

In Table 1, we have included ε for all the RBF types apart from CU and TPS. In the case of CU, including ε (i.e. using $\phi(r) = |\varepsilon r|^3$) would be pointless, since all results would turn out ε independent. Thus we always use $\phi(r) = |r|^3$ for CU. Regarding TPS, there would in fact be a weak ε dependence, which mainly would manifest itself in an additional type of singularity as $\varepsilon \to 0$, with no apparent redeeming features. Thus, all TPS results that we give are based on $\phi(r) = r^2 \log |r|$. When CU and TPS results are displayed in Figs. 9 and 10, they are thus displayed as ε-independent straight lines (although this is not strictly true in the case of TPS).

Fig. 9. The errors at $t = 10$, as functions of ε, for all the RBF choices considered in this study. In the case of GA, the ill-conditioning for RBF-Direct occurs somewhat earlier than for the other smooth RBF types, leaving a small gap (visible also in Fig. 10) between the ranges of RBF-QR and RBF-Direct. For large values of ε, the MQ results are notably more accurate than those for the other RBF types (although not nearly as good as what all the smooth RBFs achieve for small ε-values).

Fig. 10. The errors at $t = 10,000$, as functions of ε, for all the RBF choices considered in this study.

5. Analysis of the numerical results via properties of the DMs

The numerical results in Sections 3.4 and 4 can be understood through spectral analysis that is conceptually similar to Fourier analysis of equispaced finite difference (FD) and RBF approximations to the model equation

$$\frac{\partial u}{\partial t} + \frac{\partial u}{\partial x} = 0. \tag{11}$$

These steps are explained in the next subsections.

5.1. 1-D equispaced FD approximations

It has been explained and illustrated repeatedly in the literature (with [9,24] being early references) how the accuracy of a FD scheme can be understood from how it acts on individual Fourier modes. Since

$$\frac{\mathrm{d}}{\mathrm{d}x} e^{i\omega x} = i\omega e^{i\omega x},$$

$u(x) = e^{i\omega x}$ is an eigenfunction to the $\frac{\mathrm{d}}{\mathrm{d}x}$-operator. If we consider the discrete situation with a grid spacing h, the possible frequency range (due to aliasing) is $\omega \in [-\omega_{\max}, \omega_{\max}]$ where $\omega_{\max} = \frac{\pi}{h}$. Applying a FD2 (centered second order FD) scheme to $u(x) = e^{i\omega x}$ similarly gives

$$\frac{u(x+h) - u(x-h)}{2h} = \frac{e^{i\omega(x+h)} - e^{i\omega(x-h)}}{2h} = i\frac{\sin \omega h}{h} e^{i\omega x}, \tag{12}$$

i.e. the Fourier mode is again an eigenfunction, but the eigenvalue has changed from $i\omega$ to $i\frac{\sin \omega h}{h}$. Ignoring for now the factor "i", the eigenvalues are $f_{\mathrm{PS}}(\omega) = \omega$ and $f_{\mathrm{FD2}}(\omega) = \frac{\sin \omega h}{h}$, respectively. Fig. 11 illustrates these eigenvalues, as functions of ω, together with similar factors also for some FD methods of higher orders. While the PS method (limit of FD methods of increasing orders, cf. [10]) correctly treats all modes that can be represented on a grid with spacing h, lower order FD methods contain significant errors in all modes. In the context of time integration of (11), these errors correspond to errors in the phase speed of traveling waves. The

Fig. 11. Effect of $\frac{d}{dx}$ and of finite difference approximations of different orders of $\frac{d}{dx}$ when applied to a basic Fourier mode $e^{i\omega x}$. The highest mode that can be present on a grid of spacing h is denoted by $\omega_{max} = \pi/h$. (Reproduced from [10], with permission from Cambridge University Press.)

longer a time integration extends, the larger number of high modes will end up out of phase, and thereby become 'lost' as far as contributing to the overall accuracy. The lower the order of the FD approximation is, the more severe will this degradation over time become.

If we furthermore simplify by setting $h = 1$ (i.e. consider a unit-spaced infinite 1-D grid), we find

$$
\begin{aligned}
f_{FD2}(\omega) &= (\sin \omega), \\
f_{FD4}(\omega) &= (\sin \omega)\left(1 + \frac{2}{3}\left(\sin \frac{\omega}{2}\right)^2\right), \\
f_{FD6}(\omega) &= (\sin \omega)\left(1 + \frac{2}{3}\left(\sin \frac{\omega}{2}\right)^2 + \frac{8}{15}\left(\sin \frac{\omega}{2}\right)^4\right),
\end{aligned}
\tag{13}
$$

etc.

with a closed form expression for arbitrary orders given originally in [24]; see also [10, p. 41]. If we use an initial condition for (11) that, at locations $x = k$ integer, is described by

$$
u(k, 0) = \int_0^\pi \widehat{u}(\omega) \cos \omega k \, d\omega,
$$

for some function $\widehat{u}(\omega)$, the analytic-in-time FD solution at time t (again at the locations $x = k$ integer) becomes

$$
u(k, t) = \int_0^\pi \widehat{u}(\omega) \cos(\omega k - f(\omega)t) d\omega.
$$

If we also want to 'translate away' the unit-speed sideways shift of perfectly traveling waves, so that we only see the dispersive influence of the spatial approximations, we need only modify this relation to

$$
u(k, t) = \int_0^\pi \widehat{u}(\omega) \cos(\omega k + (\omega - f(\omega))t) d\omega.
\tag{14}
$$

In the case of a narrow Gaussian initial bell, direct numerical evaluation of this integral produces the solutions shown in Fig. 12. It is clear that the low order FD2 method looses its accuracy almost immediately whereas the higher order FD methods manage to keep somewhat more of the pulse integrity over longer times. However, a severe dispersive trailing wave train develops in all cases.
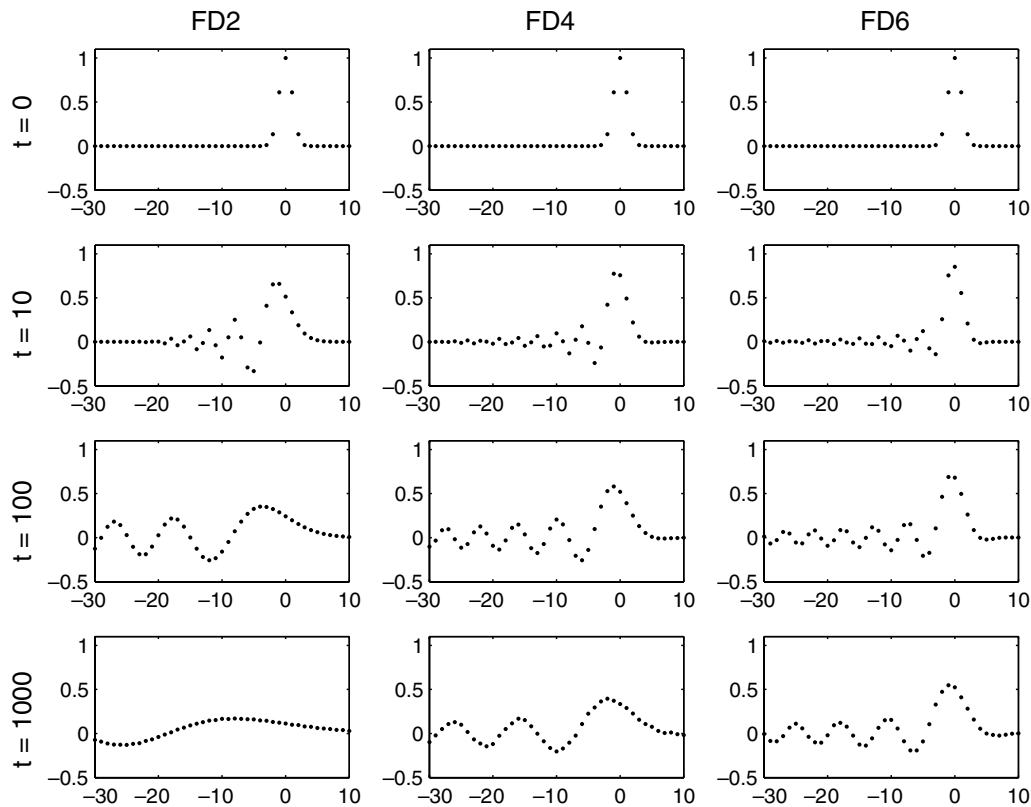
Fig. 12. The integrity of a narrow Gaussian pulse (shown as dots on a grid with $h = 1$) when integrated exactly in time, using FD2, FD4 and FD6, respectively, in space.

5.2. Dispersive errors in case of 1-D RBF approximations

Figures for RBF approximations, corresponding to Fig. 11 for FD approximations, were presented in [11]. For $\varepsilon \to 0$, the curves were seen to rapidly approach the ideal PS straight line case. In the case of IQ RBF, the counterpart to (13) was found to take a particularly simple form

$$f_{IQ}(\omega) = \omega - \frac{\pi \sinh \frac{\omega}{\varepsilon}}{\sinh \frac{\pi}{\varepsilon} \cosh \frac{\pi - \omega}{\varepsilon}}. \tag{15}$$

For the same test cases as illustrated for FD methods in Fig. 12, again by numerically evaluating (14), we obtain for IQ RBF the results seen in Fig. 13. While a large value of ε is somewhat acceptable at short times, high long-term accuracy is very clearly seen to benefit from smaller values of ε. Although the RBF scheme is spectrally accurate (as the node density is increased) for all values of ε, these solution pictures for fixed $h = 1$ and different ε are nevertheless reminiscent of the ones for increasing order FD schemes (Fig. 12). This is because of the similarities just described in how the eigenvalues of the derivative approximations vary with the Fourier frequency.

We will next carry the analysis of Sections 5.1 and 5.2 over to the case of RBF nodes scattered over the surface of a unit-sized sphere.

5.3. Dispersive errors in case of convection over the sphere

In this geometry, SPH modes $Y_\mu^\nu(\underline{x})$, $\mu \leqslant \mu_{max}$, $\nu = -\nu, \ldots, -1, 0, 1, \ldots, \nu$, form a counterpart to a truncated set of Fourier modes in a periodic 1-D case. Some low modes (up to $\mu_{max} = 4$) are illustrated in Fig. 14.

As we have already noted, when the spatial operator in (5) is approximated based on (1) or, if needed, instead by (3) or (4), all the DM's eigenvalues will be purely imaginary. From [14], we know that, as
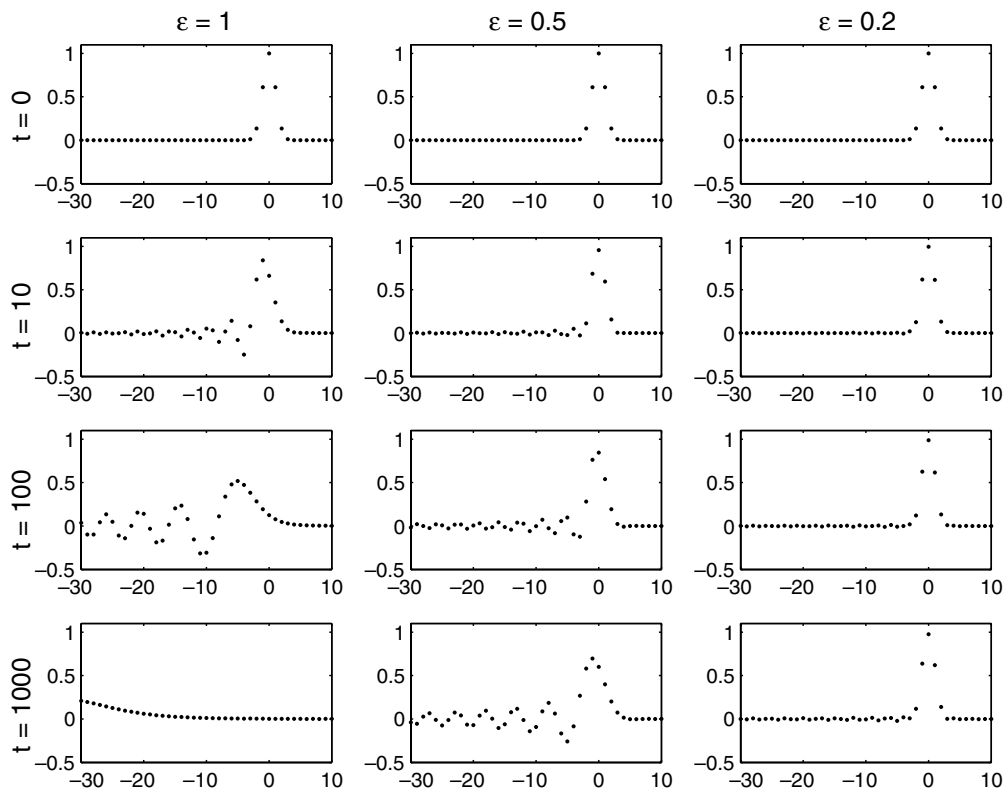
Fig. 13. Corresponding results to those in Fig. 12, but using IQ RBF with different ε-values.

Fig. 14. Illustration of the SPH basis functions Y_μ^ν for orders $\mu \leqslant \mu_{max} = 4$. The dashed lines show where they change sign.

$\varepsilon \to 0$, the space spanned by $n = \mu_{max}^2$ globally smooth RBF will exactly agree with the SPH space for $\mu \leqslant \mu_{max}$. We can now add to these observations that, in this limit of $\varepsilon \to 0$, the $n = \mu_{max}^2$ RBF DM eigenvalues will approach

$$
\begin{array}{llll}
\mu_{\max} + 1 & \text{eigenvalues} & 0 \\
\mu_{\max} & \text{eigenvalues} & +i & \text{and equally many} & -i \\
\mu_{\max} - 1 & \text{eigenvalues} & +2i & \text{and equally many} & -2i \\
\vdots \\
1 & \text{eigenvalue} & +\mu i & \text{and also one eigenvalue} & -\mu i
\end{array}
\tag{16}
$$

This can be deduced from two further observations:

(1) Truncated SPH expansions form a closed set with respect to any coordinate system rotation, making a result such as this independent on the value of α in (5).
(2) For $\alpha = 0$, an inspection of the patterns seen in Fig. 14 shows that the $\mu_{\max} + 1$ functions in the center column ($v = 0$) are unaffected by any rotation around the polar axis (leading to $\mu_{\max} + 1$ eigenvalues 0); the $\mu_{\max}$ functions for $v = +1$ and likewise for $v = -1$ repeat themselves after one revolution; the next group after 1/2 revolution, then after 1/3 revolution, etc.

Keeping $\alpha = 0$ and, to make the graphics less cluttered, choosing $\mu_{\max} = 23$ with $n = \mu_{\max}^2 = 529$ (rather than, as elsewhere this study, $n = \mu_{\max}^2 = 1849$), the eigenvalues, as functions of the SPH parameters μ and v in the case of $\varepsilon = 0$ thus become as shown graphically in Fig. 15. This flat triangular section of a plane corresponds to the PS straight line in Fig. 11 and the similar straight line for $\varepsilon = 0$ in Figs. 4.2 and 4.3 in Ref. [11]. Raising ε from zero will cause deviations from the ideal eigenvalue pattern displayed in Fig. 15. Fig. 16a displays in a different way than in Fig. 11 how the eigenvalues in the 1-D FD case vary with the order p of the FD schemes. Their extent along the imaginary axis decreases by a factor of π when we move from $p = \infty$ (PS) to $p = 2$.

Fig. 16b shows similarly how the SPH ($\varepsilon = 0$) eigenvalues, seen previously in Fig. 15, change as ε is increased from zero. In the left column, we see (according to (16)) only $2 \cdot 23 + 1 = 47$ distinct eigenvalues (out of a total of $n = 529$ eigenvalues). As ε increases, the largest eigenvalues decrease the fastest, leaving the lowest ones unchanged the longest. The 'corruption' of increasingly many eigenvalues in the RBF sphere case as ε increases is very reminiscent of how the same occurs in the 1-D equispaced case when the FD order is brought down from infinite (PS) to low order. The range of eigenvalues in these different cases will next be used to explore the time step restrictions that will need to be met for explicit time stepping methods.
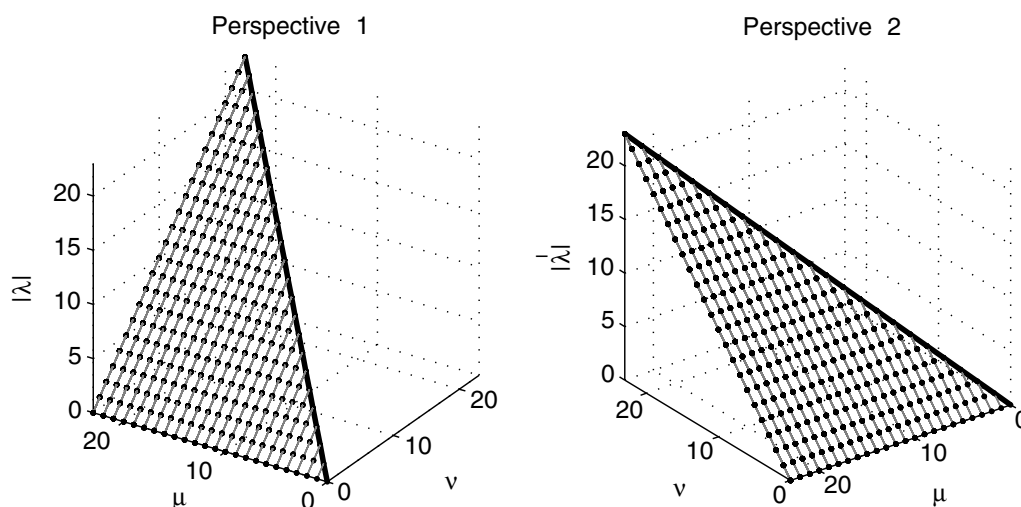
Fig. 15. The eigenvalues to the convective operator in the sphere case, computed by using IQ RBF in the $\varepsilon = 0$ limit. According to the motivation for (16), they can be associated with μ and v-values as shown.

Fig. 16. (a) The $n = 51$ 1-D convection PDE eigenvalues for different orders of FD schemes in case of $h = 1$ and a periodic domain $[-25, 25]$ (rather than $[-\infty, \infty]$). (b) The $n = 529$ sphere convection problem eigenvalues for different ε-values (in the case of MQ; results for other smooth RBF types are very similar). The vertical scales are not marked since the present issue of interest is to show qualitatively how the extents in the two cases vary with FD order and with ε-value, respectively.

5.4. Time stepping stability condition in 1-D

In the equispaced 1-D case of (11) with $\frac{\partial u}{\partial x}$ approximated by centered FD approximation of order p, the time stepping stability condition is that the eigenvalues of the spatial operator, described by $f(\omega)$ and graphically illustrated for $p = 2, 4, 6, 20, 120$ in Figs. 11 and 16a, fall within the time stepping methods stability domain, defined as usual in a complex $\xi = \lambda k$ domain, where the time stepping scheme is applied to $u_t = \lambda u$ with the time step k. For example, if we use leap-frog (LF) in time (which features as its stability domain the line section between $\xi = +i$ and $\xi = -i$), we get

FD method	Condition on k/h
FD2	< 1
FD4	< 0.7287
FD6	< 0.6305
$\cdots$	$\cdots$
PS	$< \frac{1}{\pi} \approx 0.3183$

The maximum permissible values of k/h can be read off directly from the (inverse of the) maximum values of the corresponding curves in Fig. 11, which in turn corresponds to the heights of the columns in Fig. 16a. For the LF–FD2 case, the actual stability condition thus exactly matches the result imposed by the CFL (Courant–Friedrichs–Levy) condition. As the order of spatial accuracy is increased from 2 to ∞, the actual stability condition becomes more restrictive by a factor of π (whereas the CFL condition becomes less sharp and fails to rule out increasingly long time steps).

Using (15) in place of (13), we get similar conditions on k/h (strictly valid only in our present case of $h = 1$; for details about an additional type of h-influence, see [11]):

IQ RBF	Condition on k/h
$\varepsilon = 1$	< 0.6997
$\varepsilon = 0.5$	< 0.4771
$\varepsilon = 0.2$	< 0.3803
$\cdots$	$\cdots$
$\varepsilon = 0$ (PS)	$< \frac{1}{\pi} \approx 0.3183$

The PS limit occurs in both cases (as FD order $\rightarrow \infty$ or in RBF cases as $\varepsilon \rightarrow 0$, respectively). With regard to these stability conditions, there are only very minor differences between the smooth RBF types.

5.5. Time stepping stability condition on the sphere

On the sphere, we can see from the results in Fig. 16b how the extent of the purely imaginary eigenvalue spectrum varies with ε. As we move from $\varepsilon = 0$ (SPH) to $\varepsilon \approx 1$, the spectrum shrinks by less than 10%. Importantly – and in contrast with for example spectral elements or Chebyshev-type spectral methods – there are no spurious eigenvales further out in the complex plane than where the actual physical eigenvalues are located. The estimate below is made based on $\varepsilon = 0$. Scattering $n = \mu^2$ points on the sphere, we see from (16) that the largest eigenvalue is $\mu = n^{1/2}$ in magnitude. With any type of near-uniform (locally hexagonal-like) node distribution, the distance between adjacent nodes (for $n = \mu^2$ large) becomes approximately $h = \left(\frac{8\pi}{n\sqrt{3}}\right)^{1/2}$, i.e. with LF time stepping, we would need to use

$$\frac{k}{h} < \left(\frac{\sqrt{3}}{8\pi}\right)^{1/2} \approx 0.2625. \tag{17}$$

The longest possible stable time step for any ε-value on the sphere is thus not much different from the one in the case of PS for 1-D equispaced nodes, when expressed in terms of the ratio k/h. This could have been expected since, in both test problems, the maximal physical velocity is unity and no spurious eigenvalues are present. If one uses, say, RK4 in place of LF, all the stability conditions get multiplied by $2\sqrt{2}$, i.e. for $\varepsilon = 0$ we get in 1-D $k/h < 0.9003$ and on the sphere $k/h < 0.7425$.

6. Summary and conclusions

6.1. Background to present work

The recent work by Flyer and Wright [7] clearly demonstrated some key strengths of RBF methods for solving convective-type PDEs over spherical geometries. These PDEs are particularly challenging computationally because any inaccuracies that are introduced will not only persist but will also accumulate indefinitely. At the same time, these PDEs are very important in many geophysical applications in which convection strongly dominates over dissipation. For the model problem (5), the work in [7] was limited to GA RBF with $\varepsilon \gtrsim 1$ and the only integration time that was considered was one full revolution, in our notation $t = 2\pi \approx 6.28$. In this particular case, our corresponding result can be represented as seen in Fig. 17, which can be used to fill the empty bottom right subplot space in Fig. 1. We have otherwise in this work avoided showing errors at times that correspond to an exact integer number of revolutions. This is because measuring the error as we are doing it (at node points only) would then always give zero convection errors in case of $\varepsilon = 0$, and thus might be misleadingly good also for low values of ε. However, the fact that we always start with a least square approximation means that the errors we show nevertheless are a good measure.

We have in the present work extended the earlier calculation in three aspects:

- Tested with a large number of different types of radial functions.
- Computed stably for ε-values all the way down to $\varepsilon = 0$.
- Presented results at both short ($t = 10$) and long ($t = 10,000$) times, in order to contrast time scales appropriate for weather and for climate modeling, respectively.

6.2. Main conclusions

It follows both from the present numerical results and from the accompanying analysis that very long time integrations are completely feasible with the RBF approach, but that the shape parameter ε then will need to be kept quite low, especially when integration times increase (typically bringing the computations out of reach

Fig. 17. Comparison between the results reported in [7] and the present ones for the test case of convecting the cosine bell one time around the sphere (GA RBF using $n = 1849$ ME nodes). The results are in close agreement for $\varepsilon \geqslant 3$ (where both computations were based on RBF-Direct). RBF-QR (in this case convergent only for $\varepsilon \lesssim 1$), in combination with the least squares approach, has overcome the previously seen low-ε divergence.

for the RBF-Direct approach when using standard 16 digit double precision arithmetic). All of the smooth RBFs give virtually the same accuracy, whereas the piecewise smooth ones are not at all competitive in the present context.

These general results just mentioned follow from properties of the DMs, i.e. they are not influenced by the smoothness of the convected solution. This observation is similar to the one for PS vs. FD methods, as described in Section 4.2 of [10], where it was shown that PS and high order FD methods were much better than low order FD methods, even when convecting a step function. In case of interpolation (rather than long-time advection), the situation is quite different and there will in that application probably be little point using basis functions that are much smoother than the data they are applied to.

With $n = 1849$ nodes (as used in all the calculations in the present study), it so happened (in the case of short time calculations, but not in long time ones) that RBF-Direct just barely (with the perfect choice of ε) could reach the error level that RBF-QR featured for all sufficiently small ε. The situation if n is increased further has not been explored.

It follows from the RBF-QR algorithm that $\varepsilon \to 0$ with globally smooth RBF leads to the same results as using SPH basis functions. One might therefore ask why not just use SPH as a computational basis on the sphere (as has often been done in the past). There are however several reasons in favor of an RBF approach:

- The limit $\varepsilon \to 0$ is not always the best parameter choice.
- RBFs can combine spectral accuracy with local refinement wherever this is needed (cf. discussion in [17]); SPH offer no such opportunities.
- RBF non-singularity is guaranteed whenever $\varepsilon > 0$, but not for all node sets if $\varepsilon = 0$.
- Typical SPH-PS implementations require conversions between SPH and node values on some dense grid at every time step.
- RBF codes are algebraically far less complex than SPH codes (even when implemented by means of the RBF-QR method; cf. the Appendix in [14]).

The main present concern about RBF-based methods stems from their global character, and the associated cost of full matrix×vector multiplications at each time step. Although this task is well suited for massively

parallel processing, it would nevertheless be desirable to be able to effectively apply either 'fast' (iterative) algorithms, or domain decomposition. Both these avenues have seen significant progress in recent years, but will not be surveyed here.

Appendix A. Construction of the DM with RBF-QR

The key step in the RBF-QR algorithm, as described in [14], is that a column vector containing the original RBF basis functions can be rewritten as a certain matrix product times another column vector that contains SPH functions:

$$
\begin{bmatrix} \phi(\|\underline{x} - \underline{x}_1\|) \\ \phi(\|\underline{x} - \underline{x}_2\|) \\ \vdots \\ \phi(\|\underline{x} - \underline{x}_n\|) \end{bmatrix} = [Q][E][R][\underline{y}_1(\underline{x})].
$$

Here Q is unitary, E is diagonal (with entries that are increasingly high powers of ε), and R (with more columns than rows) is upper triangular. (This equation holds to computational machine precision but not in a strict mathematical sense unless we let the number of columns in R and entries in $\underline{y}_1(\underline{x})$ be infinite.) Transposing this relation gives

$$
[\phi(\|\underline{x} - \underline{x}_1\|) \ldots \phi(\|\underline{x} - \underline{x}_n\|)] = (R\underline{y}_1(\underline{x}))^{\mathrm{T}} E^{\mathrm{T}} Q^{\mathrm{T}}. \tag{18}
$$

Next, applying this relation in turn to $\underline{x} = \underline{x}_1, \ldots, \underline{x} = \underline{x}_n$ and placing the resulting row vectors below each other allows us to write A (defined in (2) and (9)) as

$$
A = (RY_1)^{\mathrm{T}} E^{\mathrm{T}} Q^{\mathrm{T}}.
$$

All the ill-conditioning of A is now confined to the *analytically known* matrix E^{T}. Next we apply the differential operator L to (18) and obtain

$$
[L\phi(\|\underline{x} - \underline{x}_1\|) \ldots L\phi(\|\underline{x} - \underline{x}_n\|)] = (R\underline{y}_2(\underline{x}))^{\mathrm{T}} E^{\mathrm{T}} Q^{\mathrm{T}}
$$

(where $y_2(\underline{x}) = Ly_1(\underline{x})$) and therefore, after applying this to $\underline{x} = \underline{x}_1, \ldots, \underline{x} = \underline{x}_n$; cf. (8),

$$
B = (RY_2)^{\mathrm{T}} E^{\mathrm{T}} Q^{\mathrm{T}}.
$$

Because of (7), the DM then becomes

$$
D = B \cdot A^{-1} = (RY_2)^{\mathrm{T}} (E^{\mathrm{T}} Q^{\mathrm{T}}) (E^{\mathrm{T}} Q^{\mathrm{T}})^{-1} ((RY_1)^{-1})^{\mathrm{T}} = (RY_2)^{\mathrm{T}} ((RY_1)^{-1})^{\mathrm{T}} = \left((RY_1)^{-1} (RY_2) \right)^{\mathrm{T}}
$$

Computing D by means of the last expression above has avoided all the ill-conditioning even as $\varepsilon \to 0$ (eliminated through our *analytic* knowledge that $(E^{\mathrm{T}} Q^{\mathrm{T}})(E^{\mathrm{T}} Q^{\mathrm{T}})^{-1} = I$).

Appendix B. The RBF-QR method in the case of Wendland functions

Hubbert and Baxter [20] noted that RBFs, when centered at an arbitrary location $\underline{x}_i$ on the unit sphere, can be expanded in terms of SPH as

$$
\phi(\|\underline{x} - \underline{x}_i\|) = \sum_{\mu=0}^{\infty} \sum_{\nu=-\mu}^{\mu}{}' \{c_\mu Y_\mu^\nu(\underline{x}_i)\} Y_\mu^\nu(\underline{x}), \tag{19}
$$

where

$$
c_\mu = \frac{\pi}{2^{k-1}k!} \int_{-1}^{1} (1 - t^2)^k \frac{\mathrm{d}^k}{\mathrm{d}t^k} \phi\left(\sqrt{2 - 2t}\right) \mathrm{d}t. \tag{20}
$$

They also gave explicit formulas for the coefficients c_μ in several cases, such as MQ, IMQ, and GA. A formula for IQ was given in [14]. In these formulas, the potentially dangerous ε-dependence of the radial functions $\phi(\|\underline{x} - \underline{x}_i\|)$ takes the simple form of explicit high powers of ε in the coefficients, thereby permitting their analytic elimination in the RBF-QR algorithm. Typical examples include

$$\text{MQ}: c_\mu = -\varepsilon^{2\mu} \frac{2\pi(2\varepsilon^2 + 1 + (\mu + 1/2)\sqrt{1 + 4\varepsilon^2})}{(\mu + 3/2)(\mu + 1/2)(\mu - 1/2)} \left(\frac{2}{1 + \sqrt{4\varepsilon^2 + 1}}\right)^{2\mu+1},$$

$$\text{IMQ}: c_\mu = \varepsilon^{2\mu} \frac{4\pi}{(\mu + 1/2)} \left(\frac{2}{1 + \sqrt{4\varepsilon^2 + 1}}\right)^{2\mu+1}.$$

Since expansions of the type (19) have not previously been given in the case of the Wendland radial functions, and these functions' limited smoothness causes some difficulties in the RBF-QR algorithm, we give below some comments on these issues. The Wendland function W2 of order 2 can be written as

$$\phi(r) = \begin{cases} (1 - \varepsilon r)^4 (4\varepsilon r + 1) & \text{if } r < \frac{1}{\varepsilon} \\ 0 & \text{if } r \geqslant \frac{1}{\varepsilon} \end{cases}.$$

Since we are interested in these expansions (19) only for for small ε (in order to apply the RBF-QR algorithm), we assume $\varepsilon < \frac{1}{2}$. The Wendland functions on the sphere are then no longer of compact support, and the formula for W2 simplifies to

$$\phi(r) = (1 - \varepsilon r)^4 (4\varepsilon r + 1).$$

Straightforward application of (20) now gives

$$\begin{cases} c_0 = 4\pi(2 - 40\varepsilon^2 + 128\varepsilon^3 - 160\varepsilon^4 + \frac{512}{7}\varepsilon^5) \\ c_1 = \varepsilon^2 \frac{32}{63}\pi(105 - 432\varepsilon + 630\varepsilon^2 - 320\varepsilon^3) \\ c_2 = \varepsilon^3 \frac{64}{231}\pi(88 - 231\varepsilon + 160\varepsilon^2) \\ c_k = \varepsilon^3 \frac{45 \cdot 2^9 \pi((2k-5)(2k+7) - 20\varepsilon^2)}{(2k+7)(4k^2-1)(4k^2-9)(4k^2-25)} & k = 3, 4, \ldots \end{cases}$$

The expansions for W4 and W6 are easily found in a similar way, but are more lengthy and will not be explicitly given here. The fact that the expansion coefficients c_k for the smooth radial functions go to zero like $O(\varepsilon^{2k})$ is utilized in the RBF-QR algorithm. In the W2-case, c_k remains of size $O(\varepsilon^3)$ for $k = 3, 4, \ldots$, and the RBF-QR method requires a minor modification. At the key stage in that algorithm (Section 3.5 in [14]), we have now instead

$$\begin{bmatrix} \phi(\|\underline{x} - \underline{x}_1\|) \\ \phi(\|\underline{x} - \underline{x}_2\|) \\ \vdots \\ \phi(\|\underline{x} - \underline{x}_n\|) \end{bmatrix} = \begin{bmatrix} \overbrace{\frac{c_{0,\varepsilon}}{2} Y_0^0(\underline{x}_1)}^{O(1)} & \overbrace{\frac{c_{1,\varepsilon}}{1} Y_1^{-1}(\underline{x}_1) \quad \frac{c_{1,\varepsilon}}{2} Y_1^0(\underline{x}_1) \quad \frac{c_{1,\varepsilon}}{1} Y_1^1(\underline{x}_1)}^{O(\varepsilon^2)} & \overbrace{\ldots}^{O(\varepsilon^3)} \\ \frac{c_{0,\varepsilon}}{2} Y_0^0(\underline{x}_2) & \frac{c_{1,\varepsilon}}{1} Y_1^{-1}(\underline{x}_2) \quad \frac{c_{1,\varepsilon}}{2} Y_1^0(\underline{x}_2) \quad \frac{c_{1,\varepsilon}}{1} Y_1^1(\underline{x}_2) & \ldots \\ \ldots & \ldots \quad\quad \ldots \quad\quad \ldots & \ldots \\ \frac{c_{0,\varepsilon}}{2} Y_0^0(\underline{x}_n) & \frac{c_{1,\varepsilon}}{1} Y_1^{-1}(\underline{x}_n) \quad \frac{c_{1,\varepsilon}}{2} Y_1^0(\underline{x}_n) \quad \frac{c_{1,\varepsilon}}{1} Y_1^1(\underline{x}_n) & \ldots \end{bmatrix} \begin{bmatrix} Y_0^0(\underline{x}) \\ Y_1^{-1}(\underline{x}) \\ Y_1^0(\underline{x}) \\ Y_1^1(\underline{x}) \\ Y_2^{-2}(\underline{x}) \\ Y_2^{-1}(\underline{x}) \\ \vdots \end{bmatrix} = B \cdot Y$$

QR factorization of B gives

$$B = [Q] \begin{bmatrix} 1 & & & & & \\ & \varepsilon^2 & & & & \\ & & \varepsilon^2 & & & \\ & & & \varepsilon^2 & & \\ & & & & \varepsilon^3 & \\ & & & & & \ddots \end{bmatrix} \begin{bmatrix} * & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdots \\ & * & * & * & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdots \\ & & * & * & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdots \\ & & & * & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdots \\ & & & & * & * & * & * & * & * & * & * & \cdots \\ & & & & & \ddots & \cdots & \cdots & \cdots & \cdots & \cdots & \cdots & \cdots \end{bmatrix}$$

$$= Q \cdot E \cdot R.$$

After the fourth row, the elements of R do not contain factors of ε anymore, and these elements therefore no longer vanish when $\varepsilon \to 0$. Each row can be seen as a perturbed spherical harmonic function. The fact that the perturbation doesn't vanish as $\varepsilon \to 0$ means that the RBF interpolant no longer converges towards the SPH interpolant (as was the case with smooth RBFs). The situation is similar for all other non-smooth radial

functions. This also means that the convergence of the expansion will not be sped up when $\varepsilon \to 0$, but rather will be strictly algebraic in k. We will therefore need a much larger number of terms in the expansions to attain the desired accuracy in the RBF-QR method. However, since non-smooth radial functions seldom are used near their flat limit, this may be of little practical significance.

References

[1] J.P. Boyd, Chebyshev and Fourier Spectral Methods, second ed., Dover, 2001.

[2] G.L. Browning, J.J. Hack, P.N. Swarztrauber, A comparison of three different numerical methods for solving differential equations on the sphere, Mon. Wea. Rev. 117 (1989) 1058–1075.

[3] M. Buhmann, Radial Basis Functions, Cambridge University Press, Cambridge, 2003.

[4] M. Buhmann, N. Dyn, Spectral convergence of multiquadric interpolation, Proc. Edinburgh Math. Soc. 36 (1993) 319–333.

[5] T.A. Driscoll, B. Fornberg, Interpolation in the limit of increasingly flat radial basis functions, Comp. Math. Appl. 43 (2002) 413–422.

[6] G.E. Fasshauer, Meshfree Approximation Methods with Matlab, World Scientific Publishing, Singapore, 2007.

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

[8] N. Flyer, G. Wright, Solving the nonlinear shallow water wave equations using radial basis functions (submitted for publication).

[9] B. Fornberg, On a Fourier method for the integration of hyperbolic equations, SIAM J. Numer. Anal. 12 (4) (1975) 509–528.

[10] B. Fornberg, A Practical Guide to Pseudospectral Methods, Cambridge University Press, Cambridge, 1996.

[11] B. Fornberg, N. Flyer, Accuracy of radial basis function interpolation and derivative approximations on 1-D infinite grids, Adv. Comput. Math. 23 (2005) 5–20.

[12] B. Fornberg, E. Larsson, G. Wright, A new class of oscillatory radial basis functions, Comput. Math. Appl. 51 (2006) 1209–1222.

[13] B. Fornberg, D. Merrill, Comparison of finite difference and pseudospectral methods for convective flow over a sphere, Geophys. Res. Lett. 24 (1997) 3245–3248.

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

[15] B. Fornberg, G. Wright, Stable computation of multiquadric interpolants for all values of the shape parameter, Comp. Math. Appl. 48 (2004) 853–867.

[16] B. Fornberg, G. Wright, E. Larsson, Some observations regarding interpolants in the limit of flat radial basis functions, Comp. Math. Appl. 47 (2004) 37–55.

[17] B. Fornberg, J. Zuev, The Runge phenomenon and spatially variable shape parameters in RBF interpolation, Comp. Math. Appl. 54 (2007) 379–398.

[18] F.X. Giraldo, T.E. Rosmond, A scalable spectral element Eulerian atmospheric model (SEE-AM) for NWP: Dynamical core tests, Mon. Wea. Rev. 132 (2004) 133–153.

[19] J.J. Hack, R. Jakob, Description of a global shallow water model based on the spectral transform method, NCAR Technical Note TN 343 STR, 1992.

[20] S. Hubbert, B. Baxter, Radial basis functions for the sphere, in: Progress in Multivariate Approximation, International Series of Numerical Mathematics, vol. 137, Birkhauser, 2001, pp. 33–47.

[21] R. Jakob-Chien, J.J. Hack, D.L. Williamson, Spectral transform solutions to the shallow water test set, J. Comp. Phys. 119 (1995) 164–187.

[22] E.J. Kansa, Multiquadrics – a scattered data approximation scheme with applications to computational fluid-dynamics. I. Surface approximations and partial derivative estimates, Comput. Math. Appl. 19 (8/9) (1990) 127–145.

[23] E.J. Kansa, Multiquadrics – a scattered data approximation scheme with applications to computational fluid-dynamics. II. Solutions to parabolic, hyperbolic and elliptic partial differential equations, Comput. Math. Appl. 19 (8/9) (1990) 147–161.

[24] H.-O. Kreiss, J. Oliger, Comparison of accurate methods for the integration of hyperbolic equations, Tellus 24 (1972) 199–215.

[25] E. Larsson, B. Fornberg, A numerical study of radial basis function based solution methods for elliptic PDEs, Comp. Math. Appl. 46 (2003) 891–902.

[26] E. Larsson, B. Fornberg, Theoretical and computational aspects of multivariate interpolation with increasingly flat radial basis functions, Comp. Math. Appl. 49 (2005) 103–130.

[27] W.R. Madych, S.A. Nelson, Bounds on multivariate polynomials and exponential error estimates for multiquadric interpolation, J. Approx. Theory 70 (1992) 94–114.

[28] P.E. Merilees, The pseudospectral approximation applied to the shallow water equation on a sphere, Atmosphere 11 (1973) 13–20.

[29] P.E. Merilees, Numerical experiments with the pseudospectral method in spherical coordinates, Atmosphere 12 (1974) 77–96.

[30] R.D. Nair, S.J. Thomas, R.D. Loft, A discontinuous Galerkin transport scheme on the cubed sphere, Mon. Wea. Rev. 133 (2005) 814–828.

[31] R. Platte, T.A. Driscoll, Eigenvalue stability of radial basis function discretizations for time-dependent problems, Comput. Math. Appl. 51 (2006) 1251–1268.

[32] M.J.D. Powell, The theory of radial basis function approximation in 1990, in: W. Light (Ed.), Advances in Numerical Analysis, Wavelets, Subdivision Algorithms and Radial Functions, vol. II, Oxford University Press, Oxford, 1990, pp. 105–210.

[33] R. Schaback, Multivariate interpolation by polynomials and radial basis functions, Constr. Approx. 21 (2005) 293–317.

[34] R. Schaback, H. Wendland, Characterization and construction of radial basis functions, in: N. Dyn et al. (Eds.), Multivariate Approximation and Applications, Cambridge University Press, 2001, pp. 1–24.

[35] W.F. Spotz, M.A. Taylor, P.N. Swarztrauber, Fast shallow water equation solvers in latitude–longitude coordinates, J. Comput. Phys. 145 (1998) 432–444.

[36] A. St-Cyr, C. Jablanowski, J.M. Dennis, H.M. Tufo, S.J. Thomas, A comparison of two shallow water models with non-conforming adaptive grids, Mon. Wea. Rev. (in press).

[37] P.N. Swarztrauber, Spectral transform methods for solving the shallow-water equations on the sphere, Mon. Wea. Rev. 124 (1996) 730–744.

[38] M. Taylor, J. Tribbia, M. Iskandarani, The spectral element method for the shallow water equations on the sphere, J. Comput. Phys. 130 (1997) 92–108.

[39] C. Temperton, On scalar and vector transform methods for global spectral models, Mon. Wea. Rev. 119 (1991) 1303–1307.

[40] S.J. Thomas, R.D. Loft, The NCAR spectral element climate dynamical core: semi-implicit, Eulerian formulation, J. Sci. Comp. 25 (2005) 307–322.

[41] L.N. Trefethen, Spectral Methods in Matlab, SIAM, Philadelphia, 2000.

[42] H. Wendland, Scattered Data Approximation, Cambridge University Press, Cambridge, 2005.

[43] 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.

[44] R.S. Womersley, I. Sloan, Interpolation and cubature on the sphere, http://web.maths.unsw.edu.au/rsw/Sphere/.

[45] R.S. Womersley, I. Sloan, How good can polynomial interpolation on the sphere be? Adv. Comput. Math. 14 (2001) 195–226.

[46] J. Yoon, Spectral approximation orders of radial basis function interpolation on the Sobolev space, SIAM J. Math. Anal. 23 (2001) 946–958.