Boise State University

ScholarWorks

Mathematics Faculty Publications and
Presentations

Department of Mathematics

6-1-2012

A Guide to RBF-Generated Finite Differences for Nonlinear Transport: Shallow Water Simulations on a Sphere

Natasha Flyer

Institute for Mathematics Applied to Geosciences

Erik Lehto

Uppsala University

Sébastien Blaise

Universit Bordeaux

Grady Wright

Boise State University

Amik St-Cyr

Royal Dutch Shell

NOTICE: This is the author's version of a work accepted for publication by Elsevier. Changes resulting from the publishing process, including peer review, editing, corrections, structural formatting and other quality control mechanisms, may not be reflected in this document. Changes may have been made to this work since it was submitted for publication. The definitive version has been published in *Journal of Computational Physics*, 2012.
DOI: [10.1016/j.jcp.2012.01.028](https://doi.org/10.1016/j.jcp.2012.01.028)

A guide to RBF-generated finite differences for nonlinear transport: shallow water simulations on a sphere

Natasha Flyer^{a,*}, Erik Lehto^b, Sébastien Blaise^c
Grady B. Wright^{d,2}, Amik St-Cyr^{e,3}

^a*Institute for Mathematics Applied to Geosciences, National Center for
Atmospheric Research, Boulder, CO 80305 USA*

^b*Department of Information Technology, Uppsala University, SE-75105 Sweden*

^c*Institut de Mathématiques de Bordeaux, Université Bordeaux 1, 33405 Talence,
France*

^d*Department of Mathematics, Boise State University, Boise, ID 83725 USA*

^e*Royal Dutch Shell, 3737 Bellaire Boulevard, Houston, TX, 77450 USA*

.

Abstract

The current paper establishes the computational efficiency and accuracy of the RBF-FD method for large-scale geoscience modeling with comparisons to state-of-the-art methods as high-order discontinuous Galerkin and spherical harmonics, the latter using expansions with close to 300,000 bases. The test cases are demanding fluid flow problems on the sphere that exhibit numerical challenges, such as Gibbs phenomena, sharp gradients, and complex vortical dynamics with rapid energy transfer from large to small scales over short time periods. The computations were possible as well as very competitive due to the implementation of hyperviscosity on large RBF stencil sizes (corresponding roughly to 6th to 9th order methods) with up to $O(10^5)$ nodes on the sphere. The RBF-FD method scaled as $O(N)$ per time step, where N is the total number of nodes on the sphere. In the Appendix, guidelines are given on how to choose parameters when using RBF-FD to solve hyperbolic PDEs.

Key words: Radial Basis Functions, RBF, Finite Differences, RBF-FD, Hyperbolic PDEs, Spherical Geometry

1 Introduction

It has previously been shown that global radial basis functions (RBF) are highly competitive with respect to other state-of-the-art numerical methods in the arena of computational geoscience [15–18, 46]. However, they become computationally expensive when scaled to very large numbers of nodes. RBF-generated finite differences (RBF-FD) are conceptually similar to standard finite differences (FD) except the differentiation weights enforce the reproduction of RBFs rather than polynomials through the local set of nodes being considered. As in FD methods, this results in sparse matrices. However, unlike FD methods, they have all the flexibility of global RBFs in terms of handling irregular geometries and scattered node layouts. Thus, they hold the promise of computational speed and scalability to massively parallel computer platforms.

The earliest reference to RBF-FD methods seems to be a conference presentation by Tolstykh in 2000 [41]. It was independently noted (in passing only) in [13] that the convergence of RBF interpolants to polynomial form when $\varepsilon \rightarrow 0$ would suggest RBF-generated FD methods. Three independent but simultaneous works from 2003 put RBF-FD methods ‘on the map’ [33, 40, 45]. Since then they have been applied to a modest number of scientific/engineering problems such as [8–10, 32, 36] and have also been examined more conceptually in [2, 11, 12, 45, 47]. However, till now, the RBF-FD methods implemented in these studies and others have been limited to stencils with under 20 scattered nodes, corresponding to no more than a fourth-order method. Before the use of hyperviscosity in [21], larger stencil sizes could not be considered. As the stencil size increases, so does the scattering of eigenvalues of the operator matrices, resulting in loss of diagonal dominance and thus near-singular systems for elliptic operators and eigenvalues in the complex right half plane that would destabilize the time-stepping of hyperbolic systems.

The goal of this paper is to establish the capability and competitiveness of the RBF-FD method for large-scale modeling with comparisons to state-of-the-art methods and reference solutions used at National Center for Atmospheric Research (NCAR) and the Deutscher Wetterdienst (DWD - The German National Weather Service). This necessitates the use of large stencil sizes, up to

* Corresponding author.

Email addresses: flyer@ucar.edu, erik.lehto@it.uu.se, sblaise@ucar.edu, wright@math.boisestate.edu, amik.st-cyr@shell.com (Amik St-Cyr).

¹ NCAR is sponsored by the National Science Foundation. N. Flyer and E. Lehto were supported by NSF grants ATM-0620100 and DMS-0934317.

² G.B. Wright was supported by NSF grants ATM-0801309 and DMS-0934581

³ A. St-Cyr was supported by NSF grant OCI-0904599 under his former affiliation, NCAR.

101 scattered nodes, nodes sets on the sphere of $O(10^5)$, and high-order hyperviscosity of orders up to Δ^{10} . The shallow water wave equations (SWE) are used as the test bed since they capture the basic horizontal dynamics of many fluid flows. As a result, the paper is the first to test RBF-FD method for a nonlinear system of purely hyperbolic equations on a sphere and compare their performance to other high order methods on the sphere, such as discontinuous Galerkin (DG) on a cubed-sphere, spherical harmonics (NCAR/DWD-SH), and global RBFs. Evaluation of the method will be done in terms of convergence, accuracy as a function of time, and time step stability. Also, in order to assess how competitive RBF-FD are, we will perform time benchmarks against a high order discontinuous Galerkin (DG) method that was very recently developed at NCAR [4]. Since the test cases are nonlinear and highly unsteady (in one case analytically unstable), little rigorous mathematical analysis can be done and we rely on heuristic arguments to understand the results.

The paper is organized as follows: Section 2 shows how to construct the needed RBF-FD matrices; Section 3 gives the discrete representation of the SWE using RBF-FD; Section 4 discusses node layouts, hyperviscosity, and how to choose values of needed parameters in the method such as the RBF shape parameter and the order of the hyperviscosity; Section 5 considers 3 numerical studies: flow over an isolated mountain, Rossby-Haurwitz waves, and evolution of a highly nonlinear wave with complex vortical dynamics; Section 6 benchmarks the code in terms of wall-clock time and Section 7 gives the conclusions that can be drawn from the paper.

2 Constructing RBF-FD differentiation matrices

As an example, let us consider the gradient on the surface of a unit sphere in spherical coordinates (λ is longitude, θ is latitude as measured from the equator), $L = \nabla_s = \frac{1}{\cos \theta} \frac{\partial}{\partial \lambda} \hat{\lambda} + \frac{\partial}{\partial \theta} \hat{\theta}$. Assume we want to approximate L at some location x_c , denoted by the square box on any of the spheres in Figure 1(a), by a linear combination of the function values, u_k , at the neighboring n node locations, $\mathbf{x}_k$. In other words $\sum_{k=1}^n a_k u_k = (Lu)|_{x=x_c}$. The weights, a_k , are calculated by enforcing that this linear combination should be exact for RBFs, $\phi(\varepsilon \|\mathbf{x} - \mathbf{x}_k\|)_{k=1}^n$, centered at each of the node locations $\{\mathbf{x}_k\}_{k=1}^n$ (classical FD would enforce that it be exact for polynomials instead). Here, ε is a non-negative constant referred to as the shape parameter. It has also been shown through experience and studies [19, 21, 47] that better accuracy is gained by the interpolant being able to reproduce a constant. Hence, the constraint $\sum_{k=1}^n a_k = L1|_{\mathbf{x}=\mathbf{x}_c} = 0$ is added, where a_{n+1} is ignored after the

matrix is inverted. These constraints lead to the linear system of equations:

$$\begin{bmatrix} \phi(\varepsilon\|\mathbf{x}_1 - \mathbf{x}_1\|) & \phi(\varepsilon\|\mathbf{x}_1 - \mathbf{x}_2\|) & \cdots & \phi(\varepsilon\|\mathbf{x}_1 - \mathbf{x}_n\|) & 1 \\ \vdots & & \ddots & \vdots & \\ \phi(\varepsilon\|\mathbf{x}_n - \mathbf{x}_1\|) & \phi(\varepsilon\|\mathbf{x}_n - \mathbf{x}_2\|) & \cdots & \phi(\varepsilon\|\mathbf{x}_n - \mathbf{x}_n\|) & 1 \\ 1 & 1 & \cdots & 1 & 0 \end{bmatrix} \begin{bmatrix} a_1 \\ \vdots \\ a_n \\ a_{n+1} \end{bmatrix} = \begin{bmatrix} L\phi(\varepsilon\|\mathbf{x} - \mathbf{x}_1\|)|_{\mathbf{x}=\mathbf{x}_c} \\ \vdots \\ L\phi(\varepsilon\|\mathbf{x} - \mathbf{x}_n\|)|_{\mathbf{x}=\mathbf{x}_c} \\ 0 \end{bmatrix}. \quad (1)$$

If we have a total of N nodes on the sphere, then there will be N $n \times n$ matrices to invert resulting in a preprocessing cost of $O(n^3N)$. However, since $N \gg n$ in high resolution computations, the cost to time step the RBF-FD method is only $O(N)$. Its sparsity structure is seen in Figure 1b. This results in a significant speed-up from global RBFs that require $O(N^3)$ operations to create the differentiation matrices (DMs) and $O(N^2)$ to time-step. However, the method is no longer spectral and as a rough guide its order of convergence on the sphere for smooth functions is $\sqrt{n}$, e.g. a 49 point stencil is roughly a 6th order method. In our studies, N varies from 3600 to 163842 (from approximately 400km to 60km resolution, respectively) and n varies for 31 to 101 nodes.

Fig. 1. (a) An example of four 75 node RBF-FD stencils for $N = 900$ nodes on the sphere. Blue corresponds to negative differentiation weights for the convective operator and red to positive. The marker sizes reflect the magnitude of the weights. Green nodes are not included in the stencils. (b) Sparsity of the RBF-FD DM for the convective operator on a sphere using a 31 node stencil, the most commonly used in the paper, and $N = 25,600$. A kd -tree algorithm was used to order the nodes.

3 Representation of the shallow water equations

3.1 Continuous formulation

In Flyer and Wright [18], the shallow water equations on a sphere were implemented in Cartesian coordinates as follows

$$\frac{\partial u}{\partial t} = -\mathbf{p}_x \cdot \underbrace{\begin{bmatrix} (\mathbf{u} \cdot \mathbf{P}\nabla)u + f(\mathbf{x} \times \mathbf{u}) \cdot \hat{\mathbf{i}} + g(\mathbf{p}_x \cdot \nabla)h \\ (\mathbf{u} \cdot \mathbf{P}\nabla)v + f(\mathbf{x} \times \mathbf{u}) \cdot \hat{\mathbf{j}} + g(\mathbf{p}_y \cdot \nabla)h \\ (\mathbf{u} \cdot \mathbf{P}\nabla)w + f(\mathbf{x} \times \mathbf{u}) \cdot \hat{\mathbf{k}} + g(\mathbf{p}_z \cdot \nabla)h \end{bmatrix}}_{\text{RHS}}, \quad (2)$$

$$\frac{\partial v}{\partial t} = -\mathbf{p}_y \cdot \text{RHS}, \quad \frac{\partial w}{\partial t} = -\mathbf{p}_z \cdot \text{RHS}, \quad \frac{\partial h}{\partial t} = -(\mathbf{P}\nabla) \cdot (h\mathbf{u}). \quad (3)$$

Here, $\mathbf{u} = (u, v, w)$ is the wind field in directions $\mathbf{x} = (x, y, z)$, respectively, f is the Coriolis force, g is gravity and h is the geopotential height. The projection operator $\mathbf{P}$ confines the flow to the sphere and is defined as

$$\mathbf{P} = \mathbf{I} - \mathbf{x}\mathbf{x}^T = \begin{bmatrix} (1-x^2) & -xy & -xz \\ -xy & (1-y^2) & -yz \\ -xz & -yz & (1-z^2) \end{bmatrix} = \begin{bmatrix} \mathbf{p}_x^T \\ \mathbf{p}_y^T \\ \mathbf{p}_z^T \end{bmatrix}, \quad (4)$$

where $\mathbf{p}_x$ represents the projection operator in the x direction and so on. Thus, the only spatial operator that needs to be discretized is the projected gradient, $\mathbf{P}\nabla$ and its components, $\mathbf{p}_x \cdot \nabla, \mathbf{p}_y \cdot \nabla, \mathbf{p}_z \cdot \nabla$.

3.2 RBF-FD discrete formulation

Let $\phi_k(r(\mathbf{x})) = \sqrt{2(1 - \mathbf{x}^T \mathbf{x}_k)}$ be any RBF centered at $\{\mathbf{x}_k\}_{k=1}^n$, n being the number of nodes in an RBF-FD stencil. It was shown in [18] that the projected RBF gradient operator is represented as

$$\mathbf{P}\nabla\phi_k(r(\mathbf{x})) = \begin{bmatrix} x \mathbf{x}^T \mathbf{x}_k - x_k \\ y \mathbf{x}^T \mathbf{x}_k - y_k \\ z \mathbf{x}^T \mathbf{x}_k - z_k \end{bmatrix} \frac{\phi'_k(r(\mathbf{x}))}{r(\mathbf{x})} \quad (5)$$

where $'$ denotes differentiation with respect to r .

Next, we assume the function $f(\mathbf{x})$ is known at all the N node locations, $\{\mathbf{x}_j\}_{j=1}^N$, on the unit sphere. The RBF-FD stencil is a subset of those nodes of size n . As an example, let us consider $\mathbf{p}_x \cdot \nabla f(\mathbf{x})$ at a given node $\mathbf{x}_c$. We represent $f(\mathbf{x})$ as a RBF expansion of size n and apply the operator $\mathbf{p}_x \cdot \nabla$

$$[\mathbf{p}_x \cdot \nabla f(\mathbf{x})]_{|\mathbf{x}=\mathbf{x}_c} = \sum_{k=1}^n c_k \underbrace{\left[x_c \mathbf{x}_c^T \mathbf{x}_k - x_k \right]}_{b_{c,k}^x} \frac{\phi'_k(r(\mathbf{x}_c))}{r(\mathbf{x}_c)} = b^x c = (b^x A^{-1}) f = d_n^x f,$$

where A^{-1} is the inverse of the matrix defined in (1) and d_n^x are the RBF weights that will form one row of the DM, for this operator. The process is then repeated N times, via a ‘for’ loop, to form the full DM, denoted by D_N^x . For each pass through the ‘for’ loop, the n nearest neighbors to $\mathbf{x}_c$ need to be located. This can be done efficiently using a kd -tree algorithm [3], leading to the matrix structure seen in Figure 1(b). Similarly, one obtains D_N^y and D_N^z , the discrete RBF-FD approximations to the y and z components, respectively, of the projected gradient operator.

These operators are then used to construct a discrete approximation to RHS in (3):

$$\text{RHS}_D = \begin{bmatrix} \underline{u} \circ D_N^x \underline{u} + \underline{v} \circ D_N^y \underline{u} + \underline{w} \circ D_N^z \underline{u} & \begin{bmatrix} \underline{y} \circ \underline{w} - \underline{z} \circ \underline{v} \\ \underline{z} \circ \underline{u} - \underline{x} \circ \underline{w} \\ \underline{x} \circ \underline{v} - \underline{y} \circ \underline{u} \end{bmatrix} & \begin{bmatrix} D_N^x \\ D_N^y \\ D_N^z \end{bmatrix} \underline{h} \end{bmatrix},$$

where $\circ$ denotes element by element multiplication of the vectors and the underlined variables represent evaluation at the node locations. The full discretized equations are given by

$$\frac{\partial \underline{u}}{\partial t} = -\underline{p}_x \cdot \text{RHS}_D, \quad \frac{\partial \underline{v}}{\partial t} = -\underline{p}_y \cdot \text{RHS}_D, \quad \frac{\partial \underline{w}}{\partial t} = -\underline{p}_z \cdot \text{RHS}_D, \quad (6)$$

$$\frac{\partial \underline{h}}{\partial t} = \underline{u} \circ D_N^x \underline{h} + \underline{v} \circ D_N^y \underline{h} + \underline{w} \circ D_N^z \underline{h} + \underline{h} \circ (D_N^x \underline{u} + D_N^y \underline{v} + D_N^z \underline{w}), \quad (7)$$

where $\underline{p}_x, \underline{p}_y, \underline{p}_z$ are the vectors $\mathbf{p}_x, \mathbf{p}_y, \mathbf{p}_z$ evaluated at the nodes $\{\mathbf{x}_j\}_{j=1}^N$. The system is advanced in time using the classical Runge Kutta 4th-order method (RK4).

4 Nodes, parameters and hyperviscosity

This section discusses the nodes and parameters with their associated scalings that are involved in executing the RBF-FD method for hyperbolic systems on a sphere. Section 4.1 overviews the various node layouts considered. Section 4.2 discusses the need for and type of hyperviscosity used. Section 4.3 details the parameters involved in the RBF-FD method (including those for hyperviscosity) and the heuristic methodologies used to scale them.

Please refer to the Appendix for a step-by-step general guideline to choosing the parameters involved when modeling hyperbolic PDEs with RBF-FD.

4.1 Nodes layouts

We use three different kinds of node distributions on the sphere, illustrating the relative insensitivity of RBFs to the node layout that is used. The only requirement is that nodes are distinct, forming a quasi-uniform distribution over the sphere. The node sets used are minimum energy (ME) [44], maximum determinant (MD) [43, 44], and icosahedral (ICO) [1], displayed in Figure 2.

Fig. 2. The different node sets used in the numerical experiments.

4.2 Hyperviscosity

Stability issues for RBF-FD emanate from the fact that the natural intrinsic irregularity of the RBF-FD stencils causes eigenvalues of the DM to scatter into the right half of the complex plane as shown in Figure 3a. This becomes a hurdle to the RBF-FD method when 1) solving naturally dissipation-free PDEs, such that even a very mild numerical scatter of the eigenvalues into the right half complex plane can cause severe instability and 2) using large RBF-FD stencils, since as the stencil size increases so does the scatter of eigenvalues. This latter point is even an issue for systems with dissipation, in

which case the scatter is likely to be too large for the natural dissipation to control.

In this paper we have both dissipation-free PDEs and large stencils. Therefore, we need to use hyperviscosity to stabilize the RBF-FD method (to date the only known stabilization method for RBF-FD). While a regular viscosity term typically involves the Laplacian operator, a hyperviscosity operator takes the form $\gamma \cdot \Delta^k$ where γ is a very small constant, and k is some relatively large integer, typically in the range of 4 to 10, or higher still. Since the purpose of hyperviscosity is to suppress the highly oscillatory modes while leaving all the smooth ones intact, it suffices to ignore the local curvature of the sphere, and calculate $\Delta^k \phi(r)$ as if the RBF-FD stencil was located on a 2D flat plane. Even in 2D though, one can see that as k increases, the formula for calculating Δ^k can rapidly become very complicated or computationally expensive if calculated by consecutively applying the Laplacian operator. However, for Gaussian RBFs, an explicit formula is available in terms of generalized Laguerre polynomials for implementing Δ^k (see equation (3.1) and (3.2) in [21]). This renders the construction of the hyperviscosity operator both elegant and computationally inexpensive. For illustration purposes, Figures 3a and b show the effect of a Δ^4 -type hyperviscosity for $N = 400$ and $n = 31$. Although not needed for stability for the time integration lengths considered in these test cases, a filter term is added also for global RBFs, as a small increase in accuracy was noted. In this case, it is given by γA^{-1} , where A is the interpolation matrix in (1) minus the last row and column (i.e. not enforcing the constraint $\sum_{k=1}^n a_k = 0$) [21].

Fig. 3. Plots of the eigenvalue spectrum with the RK4 stability domain (solid line) for the shallow water equations linearized about the initial state of the Rossby-Haurwitz wave test case; (a) without hyperviscosity and (b) with a Δ^4 -type hyperviscosity added.

4.3 On the choice of parameters

- **RBF type:** For the global RBF method, the differences in performance between smooth RBFs are minor [23]. Inverse multiquadrics was chosen here since numerical tests suggested it to be slightly less sensitive to the choice of the shape parameter. For the RBF-FD method, Gaussian RBFs were used as this permitted the application of simple recursive formulas for the hyper-viscosity filter.
- **Node type:** All numerical results for the global method were achieved with ME nodes, as these gave marginally better results. The RBF-FD method used MD nodes since larger node sets are available. Additionally, icosahedral node sets of 40,962 and 163,842 nodes were utilized for the high resolution simulations with RBF-FD.
- **Shape parameter (ε):** To avoid saturation error (see Section 15.4 in [14]), ε was kept constant, at $\varepsilon = 5$, for the global RBF method. Numerical tests indicated that the test cases were relatively insensitive to the value of ε for $N = 3600$ to 6400 , with the later being the largest node set used for global RBFs. For RBF-FD, the node spacing may vary by orders of magnitude as the total number of nodes is increased. In order to scale to large node sets, ε is chosen such that the mean condition number of the RBF-FD interpolation matrices $\bar{\kappa}_A = \frac{1}{N} \sum_{i=1}^N (\kappa_A)_i$ is kept constant as N increases ($(\kappa_A)_i$ is the condition number of the interpolation matrix in (1), representing the i^{th} stencil). While this scaling does introduce the danger of saturation errors, convergence only for very large N , on the order of 10^5 (i.e. approximately 60km resolution), showed to be at risk. Currently there is no effective stable algorithm as $\varepsilon \rightarrow 0$ for RBF-FD on a sphere, although progress is being made in that direction. A mean condition number, $\bar{\kappa}_A$, on the order of 10^{10} to 10^{12} was found to give RBF-FD a competitive edge with regard to the accuracies reported by other high order methods in the literature.

Figures 4a and b show that for a constant mean condition number, ε varies linearly with $\sqrt{N}$. This is not surprising since the condition number strongly depends on the quantity εr , where $r \sim 1/\sqrt{N}$. Thus, to obtain a constant condition number, we let $\varepsilon(N) = c_1\sqrt{N} - c_2$, where c_1 and c_2 are constants that were fitted to the data. The values used for the numerical simulations are shown in Table 1. (For more detailed guidelines, see step 5 in Appendix).

- **Order of the hyper-viscosity (k):** The order of the Laplacian operator for the hyper-viscosity filter should be chosen large enough such that low modes are not excessively damped. However, increasing the order too much might introduce spurious eigenvalues in the filter matrix. As the stencil size n increases, an increasingly wider range of physical ‘modes’ of the PDE gets treated accurately. It is essential that the hyperviscosity filter leaves all of these basically unaffected. This is achieved by gradually

bringing up the order k of the hyperviscosity as n increases. The values in Table 1 were chosen to ensure stability while keeping diffusion as low as possible. (See step 2 in Appendix).

- **Damping coefficient (γ):** For the time integration lengths of the present test cases, the global method is stable without damping and the value $\gamma = -150 \cdot N^{-3}$ was chosen purely for optimal accuracy (negative sign due to fact we are using positive definite radial functions). For the local method, γ must be chosen with care as the method will become unstable if the damping is insufficient. With the chosen scaling of ε , it was found experimentally that the choice $\gamma = \gamma_c N^{-k}$, where k is the order of the Laplacian and γ_c is a constant, provides stability and great accuracy for all values of N considered here. This also ensures that the viscosity vanishes as $N \rightarrow \infty$. The values for γ_c are given in Table 1. Notice that the quantity $\gamma_c N^{-k}$ will vary from about $O(10^{-20})$ to $O(10^{-45})$ for the values of N considered in this paper. In other words, only a very minute amount of hyperviscosity is added to stabilize the RBF-FD method. (See step 4 in Appendix, noting also step 6.)

Table 1

The parameters for the numerical simulations.

n	$\bar{\kappa}_A$	$\varepsilon = c_1 \sqrt{N} - c_2$		$\gamma_c N^{-k} \Delta^k$	
		c_1	c_2	k	γ_c
17	$4 \cdot 10^{10}$	0.026	0.08	2	$-8 \cdot 10^{-4}$
31	$3 \cdot 10^{12}$	0.035	0.1	4	$-5 \cdot 10^{-2}$
50	$3 \cdot 10^{13}$	0.044	0.14	6	$-5 \cdot 10^{-1}$
101	$3 \cdot 10^{14}$	0.058	0.16	10	-5

(a) $n = 31$

(b) $n = 101$

Fig. 4. The mean condition number of the interpolation matrix, $\bar{\kappa}_A$, as a function of $\sqrt{N}$ and ε . a) for a stencil size of $n = 31$, b) for a stencil size of $n = 101$.

5 Numerical Studies

We will consider 3 test cases: 1) Flow over an isolated mountain that is C^0 , to study the effects of Gibbs phenomena; 2) Rossby-Haurwitz waves, to study how the method handles an unstable solution of the SWE; 3) Evolution of a highly nonlinear wave with rapid energy transfer from large to small scales, resulting in complex vortical dynamics. Since there are no analytical solutions, three high resolution reference models are considered as ‘truth’ based on, high-order DG (MUSE model), spherical harmonics (DWD model), and RBF-FD. They are described below.

- (1) DG: The DG results are computed using MUSE, the latest discontinuous Galerkin spectral element model at NCAR for the simulation of conservation laws. It has the capability of using either structured or unstructured grids combined with dynamic *hp*-adaptation to capture the multiscale aspects of the flows, efficient computational kernel implementation and high parallel performance. The model has been recently used to solve the shallow water equations on the sphere [4]. To this aim, the flows are expressed in three-dimensional Cartesian coordinates but tangentially constrained to the sphere by adding a Lagrange multiplier to the system of equations. The simulations used as references herein have been performed on a cubed sphere grid made up of 6144 elements. Each element contains 12x12 nodes to represent the solution, which results in a total of 884736 degrees of freedom and an average resolution around 26 km. For computing these reference solutions, no dissipation mechanism was found to be needed. However, for the run time versus error computations in Figure 19 in Section 6, the two dimensional exponential filter described in [28] was applied.
- (2) DWD-SH: The DWD (Deutscher Wetterdienst, German National Weather Service) spectral transform shallow water model (see <http://icon.enes.org/>) is an updated derivative of the NCAR spectral transform model. It is based on spherical harmonics (SH) implemented with de-aliasing, using Orszags 2/3 rule [27] and has become the standard reference solution in the community. For the flow over a mountain test, it has a spectral truncation of T426, that is it uses 182,329 spherical harmonic bases. For the Rossby-Haurwitz wave test, it has a spectral truncation of T511, that is it uses 262,144 spherical harmonic bases.
- (3) RBF-FD: This is a high resolution RBF-FD model based on $N = 163,824$ icosahedral nodes on the sphere, representing a 60km resolution. It uses a stencil size of $n = 31$.

5.1 Flow over an isolated mountain

This test case, originally used in [37] and popularized by [42], deals with predominantly west to east flow over a cone shaped mountain having a radius R of approximately 1925 km. The surface of the mountain is a C^0 function defined for simplicity in spherical coordinates by

$$h_{mtn} = h_{max}(1 - d/R) \quad (8)$$

where $h_{max} = 2000m$, $R = \pi/9$, and $d^2 = \min[R^2, (\lambda - \lambda_c)^2 + (\theta - \theta_c)^2]$ with (θ_c, λ_c) being the center of the mountain at latitude $30^\circ N$ and longitude $-90^\circ W$. It should be noted that in Cartesian coordinates the mountain will no longer be circular but elliptical in shape. Equation (8) is added as a forcing term to the right hand side of the PDE for the geopotential height h given in (7). Adjustment of the flow to the presence of an undifferentiable mountain results in a ringing phenomena of gravity waves whose signature echoes throughout the 15 day simulation. In order to differentiate between errors due to a non-smooth forcing, which causes Gibbs phenomena in any high-order method, and those errors inherent in the RBF-FD method, we compare the results for convergence and accuracy against test runs that use an exceptionally steep Gaussian profile that is C^∞ for the mountain, given by

$$h(\lambda, \theta) = h_{max} e^{-(2.8 \frac{d}{R})^2}, \quad (9)$$

where h_{max} , d and R are the same as for the cone mountain. The initial conditions are given by,

$$h = h_0 - \frac{1}{g}(a\Omega u_0 + \frac{u_0^2}{2})z^2, \quad \mathbf{u} = u_0\{-y, x, 0\}, \quad (10)$$

where $h_0 = 5400m$ (mean reference height), $g = 9.80616m/s^2$, $u_0 = 20m/s$, $a = 6,371,220m$ (mean radius of the earth), and $\Omega = 7.292(10)^{-5}s^{-1}$ (rotation rate of the earth). Both mountain profiles are given in Figure 5. The test case is run for 15 days. A time-series of the solution with $N = 25,600$ and $n = 31$ is given in Figure 6 for the case of the conical mountain.

a. Spatial errors and convergence

Figure 7 shows the solution at day 15 for various resolutions juxtaposed with the pointwise absolute error where the DG reference solution is used. In all cases, the stencil size is $n = 31$. With as little as $N = 3600$ nodes, the general pattern of the flow is captured with the largest error being directly downstream from the mountain. By $N = 25,600$ nodes, the error has dropped to only 2m as seen in Figure 7(f), which corresponds to a relative ℓ_∞ error of $3.4 \cdot 10^{-4}$. To better understand the convergence rate of the scheme, the normalized ℓ_2 error as a function of N is plotted in Figure 8(a) using the three different

Fig. 5. The mountain profiles with $N = 3600$.

Fig. 6. The height field h for flow over a conical mountain at day 0,5,10,15 with $N = 25,600$ and $n = 31$. The white circle indicates the base of the mountain.

reference solutions. In order to assess how RBF-FD compare to global RBF, a convergence plot for the latter is given in Figure 8(b). Note since all plots are log-log, plotting against N rather than the spatial resolution, $\sqrt{N}$, only stretches the x -axis by a factor of 2 and makes for easier reading.

Figure 8(a) shows that the normalized ℓ_2 error is an order of magnitude larger when the DWD-SH reference solution is used, as opposed to DG or the RBF-FD ($N = 163,842$) reference solutions. Furthermore when these latter reference solutions are used, the normalized ℓ_2 errors are almost identical (notice the $\circ$ overlay the $\square$ in Figure 8(a)). This same trend is also seen in Figure 8(b) with global RBFs, a different approach than both RBF-FD and DG that does not require hyperviscosity. Given that DG, RBF-FD, and global RBFs are vastly different numerical methods, strongly indicates that the DWD-SH T426

Fig. 7. Left column: The height field h for flow over a conical mountain at day 15 for varying N with $n = 31$. Contour interval is 50m. Right column: Magnitude of the error between the RBF-FD solution and the DG reference solution. Contour interval for $N=3600$ is 2m, for $N=6400$ is 1m, and for $N=25,600$ is 0.5m.

spectral simulation is providing a less accurate solution. This is further supported by the few articles that do report ℓ_2 errors for this test case [34, 35, 38], all of which use either the NCAR or DWD SH reference solution, and obtain errors on the order of 10^{-4} , which is an order of magnitude larger than that obtained by DG or RBF-FD. This is the reason that DG solution is chosen as the reference solution in Figure 7 as well as throughout the rest of the paper with regard to this test case.

Next, we consider flow over a very steep C^∞ Gaussian mountain given by Figure 5(b). Figure 9(a) displays the solution at day 15 juxtaposed with the magnitude of the error when compared to the DG reference solution. The two things to notice are that the solution looks essentially identical to the solution of the conical mountain; however, Figure 9(b) shows that there is much less

Fig. 8. The normalized ℓ_2 error in the height field h as a function of N for flow over a conical mountain at day 15 for (a) RBF-FD and (b) global RBFs. The different markers correspond to different reference solutions.

“ringing”, i.e. high frequency waves, emanating throughout the domain. In fact, if one excludes the few contours in the upper right corner of Figure 9(b), the errors (in the white region) have dropped to well below 0.1m (recalling a background mean reference height of 5400m). Another consequence of the “wringing” or Gibbs phenomena from the conical cone is that the accuracy of the RBF-FD method does not indefinitely increase with stencil size n , as shown in Figure 10(a). After $n = 31$, stencil size has no bearing on accuracy when a non-smooth mountain forcing is present. Hence, this is the stencil size chosen for this test case. In contrast, Figure 10(b) demonstrates how the accuracy of the RBF-FD method for a C^∞ solution does increase as n increases, that is as the derivative approximations become more global. However, even with a smooth forcing, the rate of convergence is not much greater than for the cone case, since both the Gaussian or cone mountains are so steep, leading to under-resolution even for very large nodes sets. To overcome this, adaptive node refinement in the area of the mountain needs to be used, as was done in [35].

(a) Day 15

(b) —Error—

Fig. 9. (a) The height field h at day 15 for flow over a Gaussian mountain with $N = 25,600$ and $n = 31$. Contour interval is 50m. (b) Magnitude of the error between the RBF-FD solution and the DG reference solution. Contour interval is 0.5m.

(a) Conical mountain

(b) Gaussian mountain

Fig. 10. The normalized ℓ_2 error in the height field h at day 15 as a function of N for different stencil sizes n for flow over a (a) conical mountain and (b) Gaussian mountain. Reference solution is the RBF-FD $N = 163,842$.

a. Temporal errors

Figure 11 shows the time trace of the normalized ℓ_2 error for the 15 day simulation using the cone mountain. As can be seen, the error grows linearly in time which is what would be expected with such a non-smooth forcing. From the end of day 1 to day 15, the normalized ℓ_2 error has increased by a factor of three. Table 5.1 gives the corresponding time steps used for each node set and the associated relative ℓ_2 error. It must be noted that these time

steps were not chosen for stability of the scheme since larger ones could have been taken, but to match spatial errors with temporal errors. This is shown in Figure 12 for both $N = 3600$ nodes and $N = 25,600$ nodes with $n = 31$. The RBF-FD spatial discretization allows for these large time steps due to the near uniform nodes. It is the temporal accuracy that actually limits the size of the time step.

Fig. 11. The normalized ℓ_2 error in the height field h as a function of time with a cone mountain forcing for $N = 25,600$ and $n = 31$.

N	Resolution (km)	Node type	Δt (minutes)	ℓ_2
3600	400	MD	20	$5.6 \cdot 10^{-4}$
6400	300		15	$2.7 \cdot 10^{-4}$
12100	220		12	$1.3 \cdot 10^{-4}$
25600	150		5	$6.0 \cdot 10^{-5}$
40962	120	ICO	3	$2.5 \cdot 10^{-5}$
163842	60		1	$1.2 \cdot 10^{-5}$

Table 2

Time steps used for the cone mountain case with respective spatial resolution and relative ℓ_2 error, using $n = 31$. The error is measured against the DG reference solution.

5.2 Rossby–Haurwitz waves

This test is proposed in [42]. The initial state is a Rossby–Haurwitz wave of wavenumber four. For the nondivergent barotropic vorticity equation, this initial condition is the solution, steadily propagating without change of shape from west to east on a sphere. However, the SWE can only approximate this motion, generating inherent instabilities in the flow [39]. As a result, this test has been controversial in its usefulness [7, 26]. Even so, it is still commonly

Fig. 12. The error at day 15 in the flow over a cone mountain test case as a function of the time step using RK4 for (a) $N = 3600$ and (b) $N = 25,600$. The error is measured against the DG reference solution.

used as a test case for novel numerical schemes in atmospheric modeling [7, 30, 35, 48]. However, it should be noted that these instabilities necessitates the use of a filter, regardless of the numerical scheme.

Figure 13 gives four RBF solutions at day 15: two RBF-FD solutions with $N = 6400$ nodes and stencil sizes of $n = 31$ and $n = 101$, Figure 13(a) and Figure 13(b), respectively, the global RBF solution with $N = 6400$, Figure 13(c), and the RBF-FD reference solution, Figure 13(d). Comparing the global RBF solution and RBF-FD reference solution to the DG and DWD-SH T511 reference solutions given in Figures 13(e) and 13(f), they are essentially indistinguishable from one another. With only $N = 6400$, RBF-FD visually captures the majority of features of the unstable flow using a stencil size of $n = 101$, which corresponds roughly to a ninth-order method. In Figure 13(a), the use of an $n = 31$ stencil shows greater differences especially at about $\pm 50^\circ$ latitude (notice concentric circular contours in blue area absent in the other figures). This may lead to the question of why an $n = 31$ stencil was used for the high resolution RBF-FD case in Figure 13(d). The reason is that for higher resolution, i.e. larger N , the results between a high and low stencil size are indistinguishable. This is reflected in Figure 14 that shows the normalized ℓ_2 error in the height field as a function of time for the 15 day run, using $N = 25,600$ nodes. Notice there is only a slight difference between using a 31 node stencil as opposed to a 101 node stencil. The likely explanation is that by this resolution we have basically resolved most of the features of the flow and increasing the order of the derivative approximations, i.e. n , one gains little in accuracy. While at coarser resolutions as $N = 6400$ one needs larger stencils to capture the higher frequencies of the flow. In other words, there is a tradeoff between the coarseness of the resolution N and order of the method needed, represented by the stencil size n . Similar tradeoffs are encountered in other contexts, such as domain decomposition schemes and ode solvers.

Next, the convergence of the RBF-FD method is considered in Figure 15 by

plotting the normalized ℓ_2 error in the height field as a function of N for the 3 different reference solutions. The line type represents the RBF-FD stencil size used, $n = 31$, $n = 101$, or global RBFs. The marker, ‘o’, ‘ $\square$ ’, or ‘ $\times$ ’, represents the reference solution used. At day 5, whether we use the RBF-FD or DG reference solution, the ℓ_2 errors for any of the cases $n = 31$, $n = 101$, or global, are almost identical (notice that the ‘ $\square$ ’ closely overlay the ‘o’). At the highest resolution of 60km ($N = 163,842$, $n = 31$), the RBF-FD method gives an error of $1.9 \cdot 10^{-5}$ when compared to the DG reference solution while with the T511 reference solution the error is about 3 times larger. Also, notice that after $N = 25,600$, the error flattens out with the T511 reference solution for the RBF-FD cases. By day 15, all the errors have increased by one to one and a half orders of magnitude. For a given case, there is not much difference in what reference solution is used, except at $N = 163,842$ where the DG reference solution gives an error that is about twice as large as the T511 solution.

Fig. 13. The solution for the Rossby-Haurwitz test case at day 15. Contour interval is 200m.

Fig. 14. A time series of the error measured against the DG reference solution for the Rossby-Haurwitz test case, with $N = 25,600$.

Fig. 15. The normalized ℓ_2 error as a function of N for the Rossby-Haurwitz test case. The different markers correspond to different reference solutions. The line type corresponds to the stencil size or global RBFs.

5.3 Evolution of highly nonlinear wave

From a numerical standpoint, this test is the most challenging in the paper, generating complex vortical dynamics [26]. It describes the evolution of a

highly nonlinear wave with rapid energy transfer from large to small scales over a short time period. The adjustment of the background flow to the initial perturbation results first in high frequency gravity waves propagating around the sphere, later followed by the roll-up of the initial flow into vortices with very sharp gradients. Due to its complexity, it has gained popularity as a test case [29, 31].

The background flow is only a function of latitude, θ , and is given by:

$$u(\theta) = \begin{cases} 0 & \text{for } \theta \leq \theta_0 \\ \frac{u_{\max}}{e_n} \exp \left[\frac{1}{(\theta - \theta_0)(\theta - \theta_1)} \right] & \text{for } \theta_0 < \theta < \theta_1 \\ 0 & \text{for } \theta \geq \theta_1 \end{cases} \quad (11)$$

where $u_{\max} = 80ms^{-1}$, $\theta_0 = \pi/7$, $\theta_1 = \pi/2 - \theta_0$, $e_n = \exp(-4(\theta_1 - \theta_0)^2)$. The initial background geopotential height field, assuming $-\pi/2 \leq \theta \leq \pi/2$, is given by

$$h(\theta) = h_0 + 1/g \int_{-\pi/2}^{\pi/2} u(\theta') [2a\Omega \cos(\theta') + \tan(\theta')u(\theta')] d\theta', \quad (12)$$

where $h_0 = 10158.295m$ (reference depth) and all other constants have the same value as in the flow over the mountain test case. To generate the instability, the height field is perturbed by Gaussians multiplied by a cosine to force the perturbation to go to zero at the poles. The perturbation $h'(\lambda, \theta)$ is given by

$$h'(\lambda, \theta) = 120 \cos(\theta) \exp \left[- \left(\frac{\lambda}{3} \right)^2 \right] \exp \left[- \left(\frac{\pi/4 - \theta}{15} \right)^2 \right] \quad \text{for } -\pi < \lambda < \pi. \quad (13)$$

Figure 16 shows the rapid evolution of the wave from day 3 to day 6 using $N = 25,600$ and $n = 101$ (before day 3 only a band of negative-blue and positive-red vorticity is seen with essentially little visible oscillations). In Figure 17, we study how well the wave is resolved as a function of both N and n . For $n = 31$, a Δ^4 - type hyperviscosity was used and for $n = 101$, a Δ^{10} - type hyperviscosity was used. For the constants in front of the hyperviscosity, refer to Table 1. There are two main concerns with this test case: a) how well the sharp gradients are resolved and b) the effect of Gibbs phenomena. As we increase both N and n , the gradients are resolved better as would be expected since both the resolution and order of the method are increasing. For the largest resolution case of $N = 163,842$ (60km or $0.54^\circ \times 0.54^\circ$), the solution is extremely similar to that given by the high-order DG solution with a resolution of 39km or $0.35^\circ \times 0.35^\circ$. However, as with regular finite differences, the higher the order of the method, i.e. n in the case, the more prominent the Gibbs phenomena. Here, $n = 101$, would correspond approximately to a

ninth-order method. Notice for the $n = 101$ plots the contour lines are more jagged, reflecting the effects of Gibbs phenomena. Another interesting feature to note is that the RBF-FD method is able to produce the basic wave pattern structure even at very coarse resolutions such as $N = 6400$ ($4.5^\circ \times 4.5^\circ$). This is not the case with a DG, spectral element method, or a finite volume method. At such coarse resolutions, the DG (MUSE) and spectral element [35] methods instead produce features of the grid such as an artificial wavenumber four pattern for the cubed-sphere. The finite volume method is noted to be so dissipative in [35] (see Figure 13) that no spatial structures appear until a resolution of $2.5^\circ \times 2.5^\circ$.

6 Time benchmarking

The RBF-FD code was benchmarked in terms of wall clock time against the DG (MUSE) code on a MacBook Pro laptop with an Intel i7 2.2GHz quad-core processor, using only a single core, and 8GB of memory. This configuration was chosen to compare the vastly different codes on the most basic level. The RBF-FD code was written in MATLAB and the DG (MUSE) code in C++. It should be mentioned that the RBF-FD method is easily parallelized on multi-core/processor machines [5, 6]. The flow over the isolated mountain test case was chosen to benchmark. The mesh specifications used in the DG timings are given in Table 6, where $N - 1$ is the highest degree polynomial used per element and R can be interpreted as how many times a face of the cubed sphere is subdivided. As can be seen in Figure 18(a), $R0$ represents the cubed sphere, for $R1$ each face of the cubed sphere has been subdivided once so that there are 4 elements per face, and so on. The RBF-FD solution of $N = 163,842$ and $n = 31$ (i.e. $\approx 60\text{km}$ resolution) was used as the reference solution in the timings. As a result, the plotted errors for the RBF-FD method in Figure 19 all use $n = 31$ with N ranging from 3136 to 40,962. For all resolutions, the RBF-FD method was computationally faster than the DG method. However, as mentioned earlier in Section 4, saturation errors in the RBF-FD method did not allow it to have a higher resolution than about 50km. This is not seen as a drawback, since saturation errors can be eliminated by doing the calculations with a smaller shape parameter ε using a flat basis function algorithm. On the sphere using global RBFs, such an algorithm has been available since 2007 [22]. For RBF-FD, algorithms are currently expected to be available within a year (personal communication Bengt Fornberg, Elisabeth Larsson, Grady Wright).

	R1N4	R1N6	R1N8	R2N6	R2N8	R3N6	R3N8
Resolution (km)	1251	834	625	417	313	208	156
Elements per face	4	4	4	16	16	64	64
Polynomial order	3	5	7	5	7	5	7
Number of d.o.f.	384	864	1536	3456	6144	13824	24576

Table 3

Samples of mesh specifications for the DG model. “d.o.f.” stands for degrees of freedom.

7 Conclusion

This paper 1) demonstrates how to implement radial basis function generated finite differences (RBF-FD) for nonlinear systems of purely convective PDEs that model complex dynamics on the sphere, specifically with regard to choices of the shape parameter and the order of the hyperviscosity with respect to the stencil size and 2) compares the performance of RBF-FD to other state-of-the-art high-order numerical methods currently used, such as a discontinuous Galerkin model recently developed at NCAR and a spherical harmonic transform model used by both NCAR and the Deutscher Wetterdienst (DWD - German National Weather Service) as a reference.

The general conclusions of the paper are:

- When modeling dynamics that have a wide range of time and spatial scales as here, the RBF-FD method requires hyperviscosity to stabilize it. However, the order of the viscosity is high, Δ^4 to Δ^{10} , depending on the stencil size, where Δ represents the Laplacian operator. This operator is then multiplied by a scaling constant such that the amount of hyperviscosity added is minute, $O(10^{-20})$ to $O(10^{-45})$.
- It is shown that the NCAR/DWD spectral transform method, based on spherical harmonics, does not seem to provide as accurate a solution as the other methods for flows where there is Gibbs phenomena, even though it has been the standard reference solution in the literature for nearly twenty years. This conclusion was based on the excellent agreement between DG, RBF-FD, and global RBFs, which gave an order of magnitude higher accuracy.
- The Rossby-Haurwitz wave test demonstrated that as the resolution increases, the difference in accuracy between using small and large stencils to calculate the derivatives at the node locations became negligible. Thus, RBF-FD can be looked at as an $O(N)$ method per time step, N being the resolution. This result was also reflected in time benchmarking.

- For the evolution of a highly nonlinear wave with sharp gradients, as the size of the stencil n increased, i.e. the order of the method increased, RBF-FD began to exhibit the expected Gibbs phenomena. However, its solution at a 60km resolution highly resembles that of the 39km resolution DG reference solution for this test. It also proved, unlike other numerical methods, to be able to capture the basic wave structure even at a very coarse resolution of 300km.
- Comparing runtimes on a laptop using a single core of a quad-core processor, the RBF-FD method was computationally faster than the DG method, from about an order and a half of magnitude for coarser resolutions to 4 times faster for the finest resolution compared. However, saturation errors in the RBF-FD method did not allow it to have a higher resolution than about 50km. That is, further increasing N requires a simultaneous increase in ϵ to maintain stability such that error stagnation results. With stable numerical algorithms that handle RBF-FD in the flat basis function regime, this need to increase ϵ with resolution is eliminated and thus so are saturation errors. Several types of stable algorithms have been developed for RBF interpolation [20, 22, 24]. Presently, projects are under way for adapting these algorithms to the task of creating general RBF-FD stencils (including approximations for hyperviscosity). It can therefore be expected that in the future saturation errors will cease to be a concern in the context of RBF-FD.

Appendix: A strategy for choosing parameters when using RBF-FD to solve hyperbolic PDEs

The steps below outline a general guideline by which we choose the parameters involved in modeling hyperbolic PDEs with RBF-FD. This appendix follows what was outlined in Sections 4.2 and 4.3. However, we do stress that some trial and error is still involved.

- (1) First, choose a stencil size n ; as n increases, so does the accuracy of the RBF-FD method. The larger n is, more eigenmodes are represented correctly.
- (2) Next, choose the order for the hyperviscosity, k , which is directly dependent on n . The larger k is, more eigenmodes are preserved correctly when the hyperviscosity is applied (see Fig. 3.1 in [21]). Table 1 shows that as n increases so does our choice for k , but by how much will depend on the PDE being solved.
- (3) For a given n and k , we choose ε for a small node set N (e.g. $N = 2500$) such that we compute at a high condition number for the RBF-FD interpolation matrix A (denoted κ_A), generally $\kappa_A = 10^8$ to 10^{12} as shown from experience. If the investigator has any result (e.g. from other model) to calibrate against, it is very useful and computationally cheap for small N to plot the error versus ε for different κ_A in order to choose which condition number works best. If not, then one should choose κ_A in approximately this range such that the answer is relatively robust to changes in $2500 \lesssim N \lesssim 5000$.
- (4) Next, for the same small N and given that we have chosen n , k , and ε from the steps above, we plot the eigenvalues of the linearized discrete right hand side operator of our PDE system. The amount the eigenvalues are pushed into the left half complex (as in Fig. 3) then becomes proportional to $\gamma_c N^{-k}$, where γ_c is a small fine tuning parameter and increases with n as shown in Table 1. The final result should be that the eigenvalues lie within the time-stepping stability domain being used but still as close as practical to the imaginary axis.
- (5) As N increases, we scale ε so that κ_A remains relatively constant. This type of scaling is shown Figure 4. Although we use the mean κ_A over all RBF-FD stencils, for roughly evenly spaced nodes, one can use any κ_A for a single stencil since they will be roughly the same. To reconstruct Figure 4 is not difficult and can be done in the following manner.

If we change $\sqrt{N} \rightarrow c\sqrt{N}$ and $\varepsilon \rightarrow c\varepsilon$, the A matrices are left invariant and therefore $\kappa_A = f(\varepsilon/\sqrt{N})$. That is κ_A is constant along straight lines

through the origin in the $(\varepsilon, \sqrt{N})$ plane. Now, it just remains to determine how κ_A depends on the slope and on n . Due to its local nature, the RBF-FD stencil on a sphere can be viewed as existing on a 2D plane. On p.389 in [25], the authors give a formula for κ_A as a function of n for the 2D non-periodic case. As an example, let $n = 31$, then $\kappa_A = c(\varepsilon/\sqrt{N})^{-14}$. Thus increasing the slope by a factor of α decreases $\log_{10} \kappa_A$ by $14 \log_{10} \alpha$. So in Figure 4(a), we only have to compute κ_A for a single case of n , N , ε and we get a good approximation of the full plot. Notice in our paper we did a small fine tuning with the ε -intercept (c_2), such that the lines do not go through the origin but around 0.1. This was only to calibrate to the DG results to demonstrate the capability of the RBF-FD method and is not necessary.

It should be noted that as we go from using node sets of $O(10^4)$ to $O(10^5)$, the κ_A we compute at increases, i.e. we switch to lines of shallower slope in Figure 4.

- (6) Lastly, we verify the stability for the larger node set using the parameter $\gamma_c N^{-k}$ computed in step 4. If unstable, increase the parameter until stability is ensured. The artificial dissipation typically requires a small safety factor to account for variability in the node sets.

References

- [1] J. Baumgardner and P. Frederickson, *Icosahedral discretization of the two-sphere*, SIAM J. Sci. Comput. **22** (1985), 1107–1115.
- [2] V. Bayona, M. Moscoso, M. Carretero, and M. Kindelan, *RBF-FD formulas and convergence properties*, J. Comp. Phys. **229** (2010), 8281–8295.
- [3] J.L. Bentley, *Multidimensional binary search trees used for associative searching*, Communications of the ACM **18** (1975), 509–517.
- [4] S. Blaise and A. St-Cyr, *A dynamic hp-adaptive discontinuous Galerkin method for shallow water flows on the sphere with application to a global tsunami simulation.*, Monthly Weather Review (2011), Accepted.
- [5] E. Bollig, N. Flyer, and G. Erlebacher, *A multi-CPU/GPU implementation of RBF-generated finite differences for PDEs on a sphere*, AGU **IN77** (2011).
- [6] ———, *Using radial basis function finite difference (RBF-FD) for PDE solutions on the GPU*, J. Comput. Phys. **submitted** (2011).
- [7] L. Bonaventura and T. Ringler, *Analysis of discrete shallow-water models on geodesic delaunay grids with c-type staggering*, Mon. Wea. Rev. **133** (2005), 2351–2373.

- [8] T. Cecil, J. Qian, and S. Osher, *Numerical methods for high dimensional Hamilton-Jacobi equations using radial basis functions*, J. Comput. Phys. **196** (2004), 327–347.
- [9] G. Chandhini and Y. V. S. S. Sanyasiraju, *Local RBF-FD solutions for steady convection-diffusion problems*, Int. J. Num. Meth. Eng. **72** (2007), 352–378.
- [10] P.P. Chinchapatnam, K. Djidjeli, P.B. Nair, and M. Tan, *A compact RBF-fd based meshless method for the incompressible navier-stokes equations*, Proc. IMechE, Part M- J. Eng. for Maritime Env. **223** (2009), 275–290.
- [11] O. Davydov and D.I. Oanh, *Adaptive meshless centres and RBF stencils for poisson equation*, J. Comp. Phy. **230** (2011), 287–304.
- [12] H. Ding, C. Shu, and D. B. Tang, *Error estimates of local multiquadric-based differential quadrature (LMQDQ) method through numerical experiments*, Int. J. Numer. Meth. Engng **63** (2005), 1513–1529.
- [13] T. A. Driscoll and B. Fornberg, *Interpolation in the limit of increasingly flat radial basis functions*, Comput. Math. Appl. **43** (2002), 413–422.
- [14] G. E. Fasshauer, *Meshless approximation methods with MATLAB*, Interdisciplinary Mathematical Sciences - Vol. 6, World Scientific Publishers, Singapore, 2007.
- [15] N. Flyer and B. Fornberg, *Radial basis functions: Developments and applications to planetary scale flows*, Computers and Fluids **26** (2011), 23–32.
- [16] N. Flyer and E. Lehto, *Rotational transport on a sphere: Local node refinement with radial basis functions*, J. Comput. Phys. **229** (2010), 1954–1969.
- [17] N. Flyer and G. B. Wright, *Transport schemes on a sphere using radial basis functions*, J. Comput. Phys. **226** (2007), 1059–1084.
- [18] ———, *A radial basis function method for the shallow water equations on a sphere*, Proc. Roy. Soc. A **465** (2009), 1949–1976.
- [19] B. Fornberg, T. A. Driscoll, G. Wright, and R. Charles, *Observations on the behavior of radial basis functions near boundaries*, Comput. Math. Appl. **43** (2002), 473–490.
- [20] B. Fornberg, E. Larsson, and N. Flyer, *Stable computations with gaussian radial functions*, SIAM J. Sci. Comput. **33(2)** (2011), 869–892.
- [21] B. Fornberg and E. Lehto, *Stabilization of RBF-generated finite difference methods for convective PDEs*, J. Comput. Phys. **230** (2011), 2270–2285.
- [22] B. Fornberg and C. Piret, *A stable algorithm for flat radial basis functions on a sphere*, SIAM J. Sci. Comput. **200** (2007), 178–192.
- [23] ———, *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.

- [24] B. Fornberg and G. Wright, *Stable computation of multiquadric interpolants for all values of the shape parameter*, Comput. Math. Appl. **48** (2004), 853–867.
- [25] B. Fornberg and J. Zuev, *The Runge phenomenon and spatially variable shape parameters in RBF interpolation*, Comput. Math. Appl. **54** (2007), 379–398.
- [26] J. Galewski, R. K. Scott, and L. M. Polvani, *An initial-value problem for testing numerical models of the global shallow-water equations*, Tellus **56A** (2004), 429–440.
- [27] D. Gottlieb and S.A. Orszag, *Numerical analysis of spectral methods: theory and applications*, SIAM Press, Philadelphia, PA, 1977.
- [28] J. S. Hesthaven, S. Gottlieb, and D. Gottlieb, *Spectral methods for time-dependent problems.*, Cambridge University Press, UK, 2007.
- [29] R.D. Nair, *Diffusion experiments with a global discontinuous Galerkin shallow-water model*, Mon. Wea. Rev. **137** (2009), 3339–3350.
- [30] R.D. Nair, S. J. Thomas, and R. D. Loft, *A discontinuous Galerkin global shallow water model*, Mon. Wea. Rev. **133** (2005), 876–888.
- [31] J.A. Pudykiewicz, *On numerical solution of the shallow water equations with chemical reactions on icosahedral geodesic grid*, J. Comput. Phys. **230** (2011), 1956–1991.
- [32] Y.Y. Shan, C. Shu, and Z.L. Lu, *Application of local MQ-DQ method to solve 3D incompressible viscous flows with curved boundary*, Comp. Modeling in Eng. & Sci. **25** (2008), 99–113.
- [33] C. Shu, H. Ding, and K.S. Yeo, *Local radial basis function-based differential quadrature method and its application to solve two-dimensional incompressible navier-stokes equations*, Comput. Meth. Appl. Mech. Engrg. **192** (2003), 941–954.
- [34] W. F. Spitz, M. A. Taylor, and P. N. Swarztrauber, *Fast shallow water equation solvers in latitude-longitude coordinates*, J. Comput. Phys. **145** (1998), 432–444.
- [35] A. St-Cyr, C. Jablonowski, J. M. Dennis, H. M. Tufo, and S. J. Thomas, *A comparison of two shallow-water models with nonconforming adaptive grids*, Mon. Wea. Rev. **136** (2008), 1898–1922.
- [36] D. Stevens, H. Power, M. Lees, and H. Morvan, *The use of PDE centers in the local RBF hermitean method for 3D convective-diffusion problems*, J. Comput. Phys. **228** (2009), 4606–4624.
- [37] L.L. Takacs, *Effects of using a posteriori methods for the conservation of integral invariants*, Mon. Wea. Rev. **116** (1988), 525–545.
- [38] M. Taylor, J. Tribbia, and M. Iskandarani, *The spectral element method for the shallow water equations on the sphere*, J. Comput. Phys. **130** (1997), 92–108.
- [39] J. Thuburn and Y. Li, *Numerical simulation of Rossby - Haurwitz waves*, Tellus **52A** (2000), 181–189.

- [40] A.I. Tolstykh, *On using radial basis functions in a “finite difference mode” with applications to elasticity problems*, Comput. Mech. **33** (2003), 68–79.
- [41] ———, *On using RBF-based differencing formulas for unstructured and mixed structured-unstructured grid calculations*, Proceedings of the 16th IMACS World Congress **228** (Lausanne (2000)), 4606–4624.
- [42] D. L. Williamson, J. B. Drake, J. J. Hack, R. Jakob, and 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.
- [43] R. S. Womersley and I. H. Sloan, *How good can polynomial interpolation on the sphere be?*, Adv. Comput. Math. **23** (2001), 195–226.
- [44] R. S. Womersley and I. H. Sloan, *Interpolation and cubature on the sphere*, Website, 2003, <http://web.maths.unsw.edu.au/~rsw/Sphere/>.
- [45] G. B. Wright, *Radial basis function interpolation: Numerical and analytical developments*, Ph.D. thesis, University of Colorado, Boulder, 2003.
- [46] G. B. Wright, N. Flyer, and D. A. Yuen, *A hybrid radial basis function - pseudospectral method for thermal convection in a 3D spherical shell*, Geophysics, Geochemistry, Geosystems **11** (2010), no. 7, Q07003.
- [47] G. B. Wright and B. Fornberg, *Scattered node compact finite difference-type formulas generated from radial basis functions*, J. Comput. Phys. **212** (2006), 99–123.
- [48] C. Yang and X.C. Cai, *Parallel multilevel methods for implicit solution of shallow water equations with nonsmooth topography on the cubed-sphere*, J. Comput. Phys. **230** (2011), 2523–2539.

(a) Day 3

(b) Day 4

(c) Day 5

(d) Day 6

Fig. 16. The evolution of the nonlinear wave. The relative vorticity for $N = 25600$, $n = 101$ is plotted. Contour interval is $2 \cdot 10^{-5} s^{-1}$ and the zero contour is not shown.

Fig. 17. The relative vorticity at day 6. Contour interval is $2 \cdot 10^{-5} s^{-1}$ and the zero contour is not shown.

Fig. 18. Mesh resolution for the DG model. The meshes used for the different simulations are obtained by splitting recursively the elements of the initial cube mesh R0.

Fig. 19. The error as a function of runtime (defined by wall clock time) for the flow over the cone mountain test case.