Published in final edited form as:

J Comput Phys. 2018 November 1; 372: 616–639. doi:10.1016/j.jcp.2018.06.036.

Hyperviscosity-Based Stabilization for Radial Basis Function-Finite Difference (RBF-FD) Discretizations of Advection-Diffusion Equations

Varun Shankar^{a,*} and Aaron L. Fogelson^b

^aDepartment of Mathematics and School of Computing, University of Utah, UT, USA

^bDepartments of Mathematics and Bioengineering, University of Utah, UT, USA

Abstract

We present a novel hyperviscosity formulation for stabilizing RBF-FD discretizations of the advection-diffusion equation. The amount of hyperviscosity is determined quasi-analytically for commonly-used explicit, implicit, and implicit-explicit (IMEX) time integrators by using a simple 1D semi-discrete Von Neumann analysis. The analysis is applied to an analytical model of spurious growth in RBF-FD solutions that uses auxiliary differential operators mimicking the undesirable properties of RBF-FD differentiation matrices. The resulting hyperviscosity formulation is a generalization of existing ones in the literature, but is free of any tuning parameters and can be computed efficiently. To further improve robustness, we introduce a simple new scaling law for polynomial-augmented RBF-FD that relates the degree of polyharmonic spline (PHS) RBFs to the degree of the appended polynomial. When used in a novel ghost node formulation in conjunction with the recently-developed overlapped RBF-FD method, the resulting method is robust and free of stagnation errors. We validate the high-order convergence rates of our method on 2D and 3D test cases over a wide range of Peclet numbers (1-1000). We then use our method to solve a 3D coupled problem motivated by models of platelet aggregation and coagulation, again demonstrating high-order convergence rates.

Keywords

Radial basis function; high-order method; hyperviscosity; meshfree; advection-diffusion

1. Introduction

Collocation methods based on radial basis functions (RBFs) are increasingly popular, due to their high-order convergence rates and their ability to naturally handle scattered node layouts on arbitrary domains. RBF interpolants can be used to generate both pseudospectral (RBF-PS) and finite-difference (RBF-FD) methods [1–6]. RBF-based methods are also easily

*Corresponding Author: vshankar@math.utah.edu (Varun Shankar), fogelson@math.utah.edu (Aaron L. Fogelson).

Publisher's Disclaimer: This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final citable form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

applied to the solution of PDEs on node sets that are not unisolvent for polynomials, such as ones lying on the sphere $\mathbb{S}^2$ [7–10] and other general surfaces [11–15].

Despite their strengths, RBF methods have historically had a major drawback: poor robustness/stability. There are two major aspects to the stability of RBF-FD methods for PDEs: *local* stability, and *global* stability. The former refers to the stability of the process of generating RBF-FD weights (interpolation), and the latter to the eigenspectrum of the sparse matrix containing the RBF-FD weights. With regards to local stability, traditional RBF interpolation (using a shape parameter) suffered from ill-conditioning, causing researchers to posit an “uncertainty principle” (a tradeoff between stability and accuracy); for a review, see [16]. Fortunately, this “local stability” problem has been solved for RBFs with shape parameters [17–22], and polyharmonic spline (PHS) RBFs [4–6, 23]. In the case of RBF-FD methods based on PHS RBFs, the key to ameliorating stability issues is the inclusion of polynomials in the approximation space and enforcing polynomial reproduction on the RBF-FD weights. Typically, however, the stencil size is required to be twice the number of polynomial basis functions [4, 23]. Thus, this technique results in larger stencils than if one were to use RBFs or polynomials alone. Fortunately, the cost associated with the increase in stencil sizes for this new *augmented RBF-FD* method can be dramatically reduced by using a variant called the *overlapped RBF-FD* method [24], developed by the first author and used in this article. Interestingly, there is also evidence that the inclusion of polynomials improves the global stability of RBF-FD methods, at least in the context of elliptic PDEs [23]. Unfortunately, in our experiments with parabolic, hyperbolic, and mixed-character PDEs, we nevertheless encountered spurious eigenvalues in RBF-FD differentiation matrices corresponding to those differential operators. Specifically, RBF-FD differentiation matrices may contain rogue eigenvalues with positive real parts, resulting in spurious growth in the numerical solution to time-dependent PDEs. In addition, when used with implicit time-stepping, we observed that the time-stepping matrices had spectral radii greater than unity.

While the theory of global stability of RBF-FD methods is virtually non-existent, Fornberg and Lehto [9] devised a practical approach to rectify the spectra of RBF-FD differentiation matrices corresponding to hyperbolic operators on the sphere by the addition of artificial hyperviscosity. They demonstrated that the technique successfully shifted any rogue eigenvalues to the left half of the complex plane. The hyperviscosity operator takes the form γ^{-k} , where $\gamma \in \mathbb{R}$ and $k \in \mathbb{N}$ are small numbers that must be selected. Then, in [10], Flyer et al. empirically computed γ and k for the shallow water equations on the sphere with two specific prescriptions: that the hyperviscosity operator vanish under refinement so the correct solution is recovered, and that the order of the operator k must be increased with the stencil size n ; this latter condition is seen in spectral methods [25]. Recently, in [4], Flyer et al. use $\gamma = (-1)^{k+1}2^{-6}h^{2k}$; however, unlike in [10], the authors of [4] do not increase k with n . In his PhD thesis [6], Barnett sets $\gamma = (-1)^{k+1}2^{-6}h^{2k+1}$, also resulting in stable simulations in Euclidean domains and on the sphere. In both these cases, 2^{-6} is a user-defined parameter whose origin is not explained. Further, and somewhat surprisingly, these expressions are independent of the magnitude of the velocity field in which the quantity of interest is advected. It is worth noting that hyperviscosity-based stabilization also has a long history of

use in spectral methods (for example, [25–27]), but in that literature, the amount of hyperviscosity is typically determined analytically.

RBF-FD methods with hyperviscosity appear to be stable despite these inconsistencies in the value of γ , possibly due to the fact that [4, 6, 9, 10] all use hand-tuned parameters in the formula for γ . Unfortunately, when we used the above recipes for hyperviscosity in the context of implicit-explicit (IMEX) linear multistep methods for solving advection-diffusion equations

$$\frac{\partial c}{\partial t} + \mathbf{u} \cdot \nabla c = \nu \Delta c, \quad (1)$$

we observed mild to severe instabilities over a range of Peclet numbers despite the presence of natural diffusion in the problem, especially on scattered and coarse node sets.

The primary contribution of this work is to carefully derive a *quasi-analytic* expression for γ , which in turn allows us to develop high-order numerical discretizations of advection-diffusion equations on scattered node sets and irregular domains. Our new expression for γ gives the correct power of h to use, an explanation for the 2^{-6} value, and the correct dependence on the velocity $\mathbf{u}$ and the diffusion coefficient ν . The resulting expression is free of any hand-tuned parameters, generalizes existing formulas in the RBF-FD literature, and connects the RBF-FD hyperviscosity literature to the spectral methods literature by which it was inspired. To derive this expression, we develop a novel semi-discrete Von Neumann analysis technique based on *auxiliary* differential operators that explicitly model the spurious growth modes introduced by RBF-FD discretizations of elliptic and hyperbolic differential operators. While we focus on advection-diffusion equations, this analysis carries over naturally to any linear PDE (or linearized nonlinear PDEs). For completeness, we also present a scaling law that relates k and the stencil size n . A secondary contribution of this work is a new scaling law connecting the power m of PHS RBFs and the degree ℓ of the appended polynomial; we find that this law significantly improves stability over fixing m as ℓ is increased, or scaling m as $2\ell + 1$. A third contribution of this work is a new ghost node method for IMEX time-stepping of advection-diffusion equations, with both diffusion and hyperviscosity being handled *implicitly* in time. While ghost nodes are not always required with RBF-FD methods [23], we found in our experiments that they help ameliorate global stability issues when used in conjunction with hyperviscosity.

The remainder of the paper is organized as follows. In Section 2, we review the overlapped RBF-FD method (with error estimates). In Section 3, we present our new scaling law for relating the PHS degree m to the polynomial degree ℓ and compare this to other possible laws. Then, in Section 4, we derive quasi-analytic expressions for the hyperviscosity parameter γ for different time integrators (explicit, implicit, and IMEX), and present efficient techniques for computing spurious growth modes; we also present an expression for k , and discuss how to approximate the hyperviscosity operator efficiently with overlapped RBF-FD. Section 5 contains our new ghost node formulation for IMEX time-stepping of advection-diffusion equations; we also briefly discuss how to deal with filling ghost values

at the initial time-step. We present 2D and 3D convergence studies in Section 6, showing results over three orders of magnitude of Peclet numbers. We apply our new techniques to solving a coupled problem inspired by mathematical models of platelet aggregation and coagulation in Section 7, and show high orders of convergence on problem in the spherical shell. We conclude with a summary and comments on future work in Section 8.

2. Overlapped RBF-FD

2.1. Mathematical description

We first present a mathematical description of the overlapped RBF-FD method, recently developed by the first author [24]. Let $X = \{x_k\}_{k=1}^N$ be a global set of nodes on the domain.

Define the stencil P_k to be the set of nodes containing node $x_{\mathcal{J}_1^k}$ and its $n-1$ nearest

neighbors $\{x_{\mathcal{J}_2^k}, \dots, x_{\mathcal{J}_n^k}\}$; here, $\{\mathcal{J}_1^k, \dots, \mathcal{J}_n^k\}$ are indices that map into the global node set X .

We defer discussion of the number of stencils to the end of this section. For the remainder of this discussion, we will focus without loss of generality on the stencil P_1 . First, define the *stencil width* ρ_1 as

$$\rho_1 = \max_j \|x_{\mathcal{J}_1^1} - x_{\mathcal{J}_j^1}\|, j = 1, \dots, n. \quad (2)$$

Choosing $\delta \in (0, 1]$, we now define the *stencil retention distance* r_1 to be

$$r_1 = (1 - \delta)\rho_1. \quad (3)$$

Here, δ is called the *overlap parameter*. Let R_1 be the set of global indices of the ρ_1 nodes in the stencil P_1 that are within the distance r_1 from $x_{\mathcal{J}_1^1}$. We write R_1 as

$$R_1 = \{\mathcal{R}_1^1, \mathcal{R}_2^1, \dots, \mathcal{R}_{\rho_1}^1\}. \quad (4)$$

In general, R_1 is some permutation of a subset of the indices of the nodes in P_1 . Let $\mathbb{B}_1$ contain the nodes whose indices are in R_1 . Thus,

$$\mathbb{B}_1 = \{x_{\mathcal{R}_1^1}, \dots, x_{\mathcal{R}_{\rho_1}^1}\}. \quad (5)$$

For convenience, we will refer to $\mathbb{B}_1$ as a *ball*, though it is merely a set of discrete points.

The overlapped RBF-FD method involves computing RBF-FD weights for all the nodes in

the ball $\mathbb{B}_1$, and repeating this process for each stencil P_k . The weights for all the nodes in $\mathbb{B}_1$ with indices in R_1 are computed using the following augmented local RBF interpolant on P_1 :

$$s_1(\mathbf{x}, \mathbf{y}) = \sum_{j=1}^n w_j^1(\mathbf{y}) \|\mathbf{x} - \mathbf{x}_{\mathcal{J}_j}^1\|^m + \sum_{i=1}^M \lambda_i^1(\mathbf{y}) \psi_i^1(\mathbf{x}), \quad (6)$$

where $\|\mathbf{x} - \mathbf{x}_{\mathcal{J}_j}^1\|^m$ is the polyharmonic spline (PHS) RBF of degree m (m is odd), and $\psi_i^1(\mathbf{x})$ are the M monomials corresponding to the total degree polynomial of degree ℓ in d dimensions. Here, the n overlapped RBF-FD weights associated with the point $\mathbf{y}$ are $w_j^1(\mathbf{y})$, $j = 1, \dots, n$. Thus, for the stencil P_1 , we have two sets of points that are relevant to the interpolant: the set P_1 itself (with $\mathbf{x} \in P_1$), and the set $\mathbb{B}_1$ (with $\mathbf{y} \in \mathbb{B}_1$). In order to compute the weights for the linear operator $\mathcal{L}$ uniquely at all nodes in $\mathbb{B}_1$ with indices in the set R_1 , we impose the following two (sets of) conditions:

$$s_1|_{\mathbf{x} \in P_1, \mathbf{y} \in \mathbb{B}_1} = \mathcal{L} \|\mathbf{x} - \mathbf{x}_{\mathcal{J}_j}^1\|^m|_{\mathbf{x} \in \mathbb{B}_1}, j = 1, \dots, n, \quad (7)$$

$$\sum_{j=1}^n w_j^1(\mathbf{y}) \psi_i^1(\mathbf{x})|_{\mathbf{x} \in P_1, \mathbf{y} \in \mathbb{B}_1} = \mathcal{L} \psi_i^1(\mathbf{x})|_{\mathbf{x} \in \mathbb{B}_1}, i = 1, \dots, M. \quad (8)$$

The first set of conditions enforces that $s_1(\mathbf{x}, \mathbf{y})$ interpolate the derivatives of the PHS RBF at all nodes in $\mathbb{B}_1$. The second set of conditions enforces polynomial reproduction on the overlapped RBF-FD weights. If ℓ is the degree of the appended polynomial $\psi^1(\mathbf{x})$, $M = \binom{\ell + d}{d}$; for stability, we also require that $M \leq \lfloor \frac{n}{2} \rfloor$ [4, 5, 24]. The interpolant (6) and the two conditions (7)–(8) can be collected into the following block linear system:

$$\begin{bmatrix} A_1 & \Psi_1 \\ \Psi_1^T & O \end{bmatrix} \begin{bmatrix} W_1 \\ \Lambda_1 \end{bmatrix} = \begin{bmatrix} B_{A_1} \\ B_{\Psi_1} \end{bmatrix}, \quad (9)$$

where

$$(A_1)_{ij} = \|\mathbf{x}_{\mathcal{J}_i}^1 - \mathbf{x}_{\mathcal{J}_j}^1\|^m, i, j = 1, \dots, n, \quad (10)$$

$$(\Psi_1)_{ij} = \psi_j^1(\mathbf{x}_{\mathcal{J}_i^1}); i = 1, \dots, n, j = 1, \dots, M, \quad (11)$$

$$(B_{A_1})_{ij} = \mathcal{L} \|\mathbf{x} - \mathbf{x}_{\mathcal{R}_j^1}\|^m |_{\mathbf{x} = \mathbf{x}_{\mathcal{J}_i^1}}, i = 1, \dots, n, j = 1, \dots, p_1, \quad (12)$$

$$(B_{\Psi_1})_{ij} = \mathcal{L} \psi_i^1(\mathbf{x}) |_{\mathbf{x} = \mathbf{x}_{\mathcal{R}_j^1}}, i = 1, \dots, M, j = 1, \dots, p_1, \quad (13)$$

$$O_{ij} = 0, i, j = 1, \dots, M. \quad (14)$$

W_1 is the *matrix* of overlapped RBF-FD weights, with each column containing the RBF-FD weights for a point $\mathbf{x} \in \mathbb{B}_1$. The linear system (9) has a unique solution if the nodes in P_1 are distinct [16, 28]. The matrix of polynomial coefficients Λ_1 is merely a set of Lagrange multipliers that enforces the polynomial reproduction constraint (8).

2.2. Algorithmic description of differentiation matrix assembly

We next describe how to assemble an $N \times N$ sparse differentiation matrix from the mathematical description above. The algorithm described here is a *serial* one; parallel algorithms for assembly are a subject of ongoing research. The procedure involves using a kd-tree for nearest neighbor searches, together with a form of breadth-first search, and is described in Algorithm 1. For simplicity, algorithm 1 describes the assembly process as directly populating the entries of L . In practice, one stores the non-zero elements, and row and column indices/pointers (depending on the sparse matrix format in use), and uses these arrays in place of L .

It is also worth briefly examining the cost of this process. Since Algorithm 1 uses a kd-tree, it has a preprocessing cost of $O(N \log N)$. Each iteration in the outer loop of the algorithm (line 4) incurs the following costs: $O((n-1) \log N)$ to find the nearest neighbors, and $O((n+M)^3 + p_k(n+M)^2)$ for finding the weight matrix W_k . There are N_δ such steps. If the points are quasi-uniform, then $N_\delta \approx \frac{N}{p}$, where $p = (1-\delta)^{d_n}$ [24]. If $\delta < 1$, each iteration of this loop is more expensive than the standard RBF-FD method, but there are far fewer iterations (by a factor $\frac{1}{p}$). The assembly process as described above also requires fewer nearest neighbor searches than the standard RBF-FD method if $\delta < 1$.

2.3. Local error estimates

The polynomial reproduction property (8) of the augmented local RBF interpolant can be used to develop a local error estimate [29]. Following [24], if $\mathcal{L}$ is a differential operator of order θ applied to some function f , this error estimate can be written as:

$$\left| \mathcal{L}f(\mathbf{y}) - \sum_{j=1}^n w_j(\mathbf{y})f(\mathbf{x}_j) \right| \leq P(\mathbf{y}) \max_{|\alpha|=\ell} \|D^\alpha f\|_{W^{1,\infty}(\Omega_1)} (h(\mathbf{y}))^{\ell+1-\theta}, \quad (15)$$

Algorithm 1

Differentiation matrix assembly using Overlapped RBF-FD

Given: $X = \{\mathbf{x}_k\}_{k=1}^N$, the set of nodes in the domain.

Given: $\delta \in (0, 1]$, the overlap parameter.

Given: $\mathcal{L}$, the linear differential operator to be approximated.

Given: $n \ll N$, the stencil size.

Generate: L , the $N \times N$ differentiation matrix approximating $\mathcal{L}$ on the set X .

Generate: N_δ , the number of stencils.

- 1: Build a kd-tree on the set X in $O(N \log N)$ operations.
- 2: Initialize g , an N -long array of flags, to 0.
- 3: Initialize the stencil counter, $N_\delta = 0$.
- 4: **for** $k = 1, N$ **do**
- 5: **if** $g(k) == 0$ **then**
- 6: Use kd-tree to get $\{\mathbf{x}_{\mathcal{I}_1^k}, \dots, \mathbf{x}_{\mathcal{I}_n^k}\}$. Here, $\mathcal{I}_1^k = k$.
- 7: Get $R_k = \{\mathcal{R}_1^k, \dots, \mathcal{R}_{p_k}^k\}$ and $\mathbb{B}_k = \{\mathbf{x}_{\mathcal{R}_1^k}, \dots, \mathbf{x}_{\mathcal{R}_{p_k}^k}\}$ using (2)-(5).
- 8: Use (9) to compute W_k , the $n \times p_k$ matrix of RBF-FD weights.
- 9: **for** $i = 1, p_k$ **do**
- 10: Set $g(\mathcal{R}_i^k) = 1$.
- 11: **for** $j = 1, n$ **do**
- 12: Set $L(\mathcal{R}_i^k, \mathcal{I}_j^k) = W_k(j, i)$.
- 13: **end for**
- 14: **end for**
- 15: Set $N_\delta = N_\delta + 1$.
- 16: **end if**
- 17: **end for**

where $\mathbf{y}$ is some evaluation point, $h(\mathbf{y})$ is the largest distance between the point $\mathbf{y}$ and every point $\mathbf{x}$ in the stencils, $P(\mathbf{y})$ is some *growth function* [29], and α is a multiindex. The term $\max_{|\alpha|=\ell} \|D^\alpha f\|_{W^{1,\infty}(\Omega_1)}$ is simply the Sobolev ∞ -norm of f on the stencil P_1 with convex

hull Ω_1 . The error estimate shows that the order of convergence in any RBF-FD method depends primarily on the degree ℓ not the location of the evaluation point $\mathbf{y}$. One can therefore expect the same orders of convergence for the overlapped RBF-FD method as in the standard one, assuming $\mathbf{y}$ is not too close to the boundary of Ω_1 . As was shown in [24], the difference in accuracy between the standard and overlapped RBF-FD methods is virtually non-existent for higher-order approximations, can be kept small for lower-order approximations using judicious choices of δ .

2.4. Parameter selection

We now describe how the different parameters in the overlapped RBF-FD method are selected. Based on (15), we know the polynomial degree ℓ controls the order of the approximation for a differential operator of order θ . Thus, given a differential operator of order θ , if we require an RBF-FD method with order of accuracy ξ , we set

$$\ell = \xi + \theta - 1. \quad (16)$$

For instance, if we are approximating the gradient operator ∇ ($\theta = 1$) to third-order accuracy ($\xi = 3$), this gives $\ell = 3$. On the other hand, if we are approximating the Laplacian operator ($\theta = 2$) to third-order accuracy ($\xi = 3$), this gives $\ell = 4$. The stencil size n in d dimensions can be deduced from the relationships $M = \binom{\ell + d}{d}$ and $M \leq \lfloor \frac{n}{2} \rfloor$ to be $n \geq 2M$. In practice, it appears to be beneficial to use a few more nodes per stencil than $2M$, though the optimal choice is likely problem dependent. We use

$$n = 2M + \lfloor \ln(2M) \rfloor. \quad (17)$$

We must also choose the overlap parameter $\delta \in (0, 1]$. For small values of n (equivalently of ℓ), this choice trades off speedup with accuracy as $\delta \rightarrow 0$; however, as n increases, smaller values of $\delta > 0.2$ barely impact accuracy [24]. In practice, we have observed that setting $\delta = 0.2$ typically completely decouples the stencils, resulting in ill-posed subproblems when solving PDEs. Given these constraints, we use the following heuristic:

$$\delta = \begin{cases} 0.7 & \text{if } \ell \leq 3 \\ 0.5 & \text{if } 4 \leq \ell < 6 \\ 0.3 & \text{if } \ell \geq 6 \end{cases}$$

The automatic selection of δ based on stability properties of augmented local RBF interpolants is a subject of ongoing research. It is likely that the optimal choice of δ depends on the local Lebesgue functions corresponding to the differential operator $\mathcal{L}$; see Section 3 and [24, 29] for discussions on these local Lebesgue functions.

Connection to other RBF methods—As described here, the overlapped RBF-FD method bears some resemblance to both the standard RBF-FD method and the RBF-Partition of Unity (RBF-PU) method [16, 28]. In the latter case, one computes a *global*

interpolant to the node set X by dividing X into a set of overlapping partitions called *patches* (controlled by an overlap parameter), forming a local RBF approximation on each patch, and then blending these approximations together with weight functions that are compactly-supported on each patch. As such, the standard RBF-FD method can be thought of as a very special RBF-PU method with weight functions that are exactly 1 at each node, and zero elsewhere, resulting in as many patches as there are nodes. The overlapped RBF-FD method can then be thought of as an RBF-PU method where the weight functions are 1 on the balls $\mathbb{B}_k$, zero elsewhere, resulting in fewer patches than nodes. However, in our opinions, the overlapped RBF-FD method is best viewed as an RBF-FD method where one makes a choice of either centered or one-sided approximations for each node in X , thereby resulting in stencil-sharing for some nodes. Indeed, the overlapped and standard RBF-FD methods share the same local error estimates from Section 2.3. That said, it may be that viewing the assembly process from the RBF-PU perspective may help with parallelizing our method. This is a topic of ongoing research.

3. A new scaling law for m and ℓ

The selection of the degree m of the PHS RBF requires more thought. On the one hand, m is known to control the constant in the RBF-FD error estimate (15), with higher values of m giving lower errors [5, 6]. On the other hand, larger values of m may cause greater instability, as the smallest eigenvalue of the block matrix from (9) can be bounded from below by q^m , where q is the separation distance of the stencil P_i ; see [28] for some preliminary estimates, and [4] for remarks about stability. A third more subtle issue is the question of whether m and ℓ should be related. Historically, it was common to use the *classical scaling law* [30]:

$$m = 2\ell + 1, \quad (18)$$

where ℓ is selected to be the minimum to prove unisolvency for the RBF interpolant with $\phi(r) = r^m$. This approach was also recently used for PHS-based methods for advection on the sphere [31], and advection and reaction-diffusion equations on more general manifolds [32]. However, modern PHS RBF-FD methods for Euclidean domains decouple m and ℓ [4–6, 24]. Our goal is to attempt to find a scaling law that produces the smallest spurious eigenvalue (if any) in the spectrum of the discrete Laplacian L . To that end, we explored two alternatives to the classical scaling law. The first alternative is:

$$m = \begin{cases} \ell & \text{if } \ell \text{ is odd} \\ \ell + 1 & \text{if } \ell \text{ is even} \end{cases} \quad (19)$$

which serves to gently increase m as ℓ is increased. The second alternative, which we ultimately settled on, is:

$$m = \begin{cases} \ell & \text{if } \ell \text{ is odd} \\ \ell - 1 & \text{if } \ell \text{ is even} \end{cases} \quad (20)$$

The eigenvalues for the discrete Laplacian L computed using the classical scaling law, these two alternatives, and the choice of $m = 3$ are shown for $N = 658$ nodes on the disk using $\ell = 6$ in Figure 1. From Figure 1a, it is clear that the classical scaling law $m = 2\ell + 1$ results in a large spurious eigenvalue. In general, all the other scaling laws appear to give very similar results, albeit with some subtle differences. Figure 1b uses the scaling law $m = \ell + 1$, while Figure 1c uses the scaling law $m = \ell - 1$; clearly, both produce almost identical results, with the exception that $m = \ell + 1$ results in a slightly larger spurious eigenvalue. Continuing this trend, it may be tempting to always simply choose a very small reasonable value of m . However, Figure 1d with $m = 3$ shows that this may not be entirely advisable either, since the largest spurious eigenvalue here is actually larger than the cases of $m = \ell + 1$ and $m = \ell - 1$. The spectrum of L for different scaling laws for $N = 919$ nodes with $\ell = 6$ is shown in Figure 2. Once again, $m = 2\ell + 1$ can be rejected as extremely unstable, even after refinement (Figure 2a). However, the gaps between the other scaling laws have widened slightly as well. For instance, $m = \ell + 1$ (Figure 2b) now results in a spurious eigenvalue that is twice as large as $m = \ell - 1$ (Figure 2c). Even more interestingly, $m = 3$ (Figure 2d) results in a spurious eigenvalue whose real part is much larger than the $m = \ell + 1$ and $m = \ell - 1$ cases. Once again, the most stable choice appears to be $m = \ell - 1$. Picking a fixed value of $m = 3$ (or any other number) regardless of the choice of ℓ may be inadvisable on Euclidean domains except when the node sets are very regular. These trends continue for different values of ℓ (not shown). It may be the case that the correct scaling relationship needs to be reevaluated for each type of node set used.

Thus far, we have only discussed the *global* stability associated with the different scaling laws. It is reasonable to expect that these global patterns are reflected in some way at the stencil level. In other words, we expect that the global stability must be connected to *local* stability in some way. Before we can explore this connection, we must first define a few quantities. Returning to the stencil P_1 , consider the *local $\mathcal{L}$ -Lebesgue function* of the augmented RBF interpolant $s_1(\mathbf{x}, \mathbf{y})$ used for approximating the operator $\mathcal{L}$, defined as

$$\Lambda_{\mathcal{L}}^1(\mathbf{y}) = \sum_{j=1}^n |w_j^1(\mathbf{y})| = \| \mathbf{W}^1 \|_1, \quad (21)$$

that is, the 1-norm of the RBF-FD weight vector at a point. Let $w_1^1(\mathbf{y})$ be the RBF-FD weight associated with the stencil point $\mathbf{y}$. From [24], we have the sufficient condition for all eigenvalues of the differentiation matrix L (corresponding to $\mathcal{L}$) having non-positive real parts as:

$$2w_1^1(\mathbf{y}) \geq -\Lambda_{\mathcal{L}}^1(\mathbf{y}). \quad (22)$$

This condition highlights the connection between the local Lebesgue function (itself connected to the stability of the local RBF interpolant [28]) and the *global* stability of the RBF-FD method. In general, large values of $\Lambda_1^{\mathcal{L}}(y)$ are associated with instability, and typically produce spurious eigenvalues with positive real parts [24]. To examine the impact of the scaling laws (18)–(20) on the local Lebesgue functions (and therefore local stability), we will now study the local Δ -Lebesgue functions for a set of $N=4000$ Halton points in $[-1, 1]^2$, where Δ is the Laplacian; evenly-spaced points are used on the boundary; these Halton points are typically even more scattered than Poisson disk samples, and represent a worst-case scenario. The results are shown in Figure 3. The first column in Figure 3 corresponds to the classical law (18), the second to (19), and the third to (20). The colorbar scale is fixed across each row, but allowed to increase between rows to make clear that RBF-FD weights can increase in magnitude as the polynomial degree is increased. Immediately, it is clear that the classical scaling law results in the largest weights, and is consequently almost always unstable (Figures 3a,3d, 3g). Careful observation also shows that the last column (corresponding to (20)) results in more stable approximations than the middle column (corresponding to (19)), which matches our conclusions from the eigenvalue plots (Figures 1 and 2). In general, we find that the scaling law (20) produces stable approximations across a wide range of parameters, and therefore we use it in this article.

4. Hyperviscosity-based Stabilization

As mentioned previously, the RBF-FD differentiation matrices corresponding to the discretized gradient operator (and occasionally even the Laplacian operator) can contain eigenvalues with positive real parts. Such eigenvalues can cause spurious growth in the numerical solutions of advection-diffusion equations. Our approach to ameliorating this issue is to add a small vanishing amount of artificial hyperviscosity to the model. This transforms Equation 1 to

$$\frac{\partial c}{\partial t} = -\mathbf{u} \cdot \nabla c + \nu \Delta c + \gamma \Delta^k c, \quad (23)$$

where $\gamma \Delta^k c$ is the artificial hyperviscosity term. This approach has been used in conjunction with explicit time-stepping in convective PDEs [4, 6, 9, 10]. Our goals in this section are twofold: first, to present expressions for γ and k , and second, to discuss the spatial approximation of the hyperviscosity term in the context of the RBF-FD method. We proceed by first deriving expressions for γ in the advection-dominated regime, in the diffusion-dominated regime, then for the case where the two processes co-exist.

4.1. Selecting γ in the advection-dominated regime ($\nu \rightarrow 0$)

As mentioned in Section 1, the literature on modern RBF-FD methods gives a prescription for γ . In [4, 6], for instance, Flyer et al. set this parameter to $\gamma = (-1)^{1-k} 2^{-6} h^{2k-1}$, where k is the power of the Laplacian. We will show with our analysis that this formula is a special case of a much more general formula.

Assume for now a simplified 1D version of (1) where $v = 0$. This reduces it to the advection equation

$$\frac{\partial c}{\partial t} + u \frac{\partial c}{\partial x} = 0. \quad (24)$$

Assume that c is a plane wave of the form $c(x, t) = \hat{c}(t)e^{ikx}$, where $\hat{k}$ is the wave number. The above equation should simply translate the plane wave to the right at a speed u .

Unfortunately, the problem lies in the approximation of $\frac{\partial}{\partial x}$. In general, an RBF-FD discretization will instead result in our solving the *auxiliary* equation

$$\frac{\partial c}{\partial t} + u \frac{\tilde{\partial} c}{\partial x} = 0, \quad (25)$$

where $\frac{\tilde{\partial} c}{\partial x}$ is an *auxiliary* differential operator that introduces a new term $\varepsilon \hat{k}^q$ into the right hand side of the above equation so that we have

$$\frac{\partial c}{\partial t} + u i \hat{k} \hat{c}(t) e^{ikx} = u \varepsilon \hat{k}^q \hat{c}(t) e^{ikx}, \quad (26)$$

where $\varepsilon = 0$ causes *growth* in c , and $q \in \mathbb{R}$. In general, ε and q are unknown *a priori*, but are functions of the node set and basis functions used; see [34] for a discussion on the relationships between eigenvalues of differential operators, RBFs, and node sets. We assume in this work that an upper-bound for ε is the real part of the spurious eigenvalue with the largest real part. In the RBF-FD context, we wish to add some amount of artificial hyperviscosity to stabilize the auxiliary PDE, giving us

$$\frac{\partial c}{\partial t} + u \frac{\tilde{\partial} c}{\partial x} = \gamma \frac{\partial^{2k} c}{\partial x^{2k}}, \quad k \in \mathbb{N}. \quad (27)$$

We will now derive expressions for γ that *cancel* out the spurious growth mode $\varepsilon \hat{k}^q$.

4.1.1. Time-stepping with explicit RK methods—Our goal in this section is to derive a hyperviscosity formulation for explicit RK methods. We first examine the forward Euler discretization as this is instructive. Discretizing (27) using the Forward Euler scheme, we obtain

$$c^{n+1} - c^n + \Delta t u \frac{\tilde{\partial} c^n}{\partial x} = \Delta t \gamma \frac{\partial^{2k} c^n}{\partial x^{2k}}, \quad (28)$$

where the superscripts indicate time levels. Now, akin to Von Neumann analysis, assume that $c^{n+1} = Gc^n$, where G is some growth factor. For time stability, we require that $|G| \leq 1$. Substituting the plane wave definition of c^n , we obtain

$$G\hat{c}(t_n)e^{ikx} = \hat{c}(t_n)e^{ikx} - \Delta t u \hat{c}(t_n)e^{ikx}(ik - \epsilon \hat{k}^q) + \Delta t \gamma i^{2k} \hat{k}^{2k} \hat{c}(t_n)e^{ikx}. \quad (29)$$

This simplifies to

$$G = (1 - \Delta t u i \hat{k}) + (u \Delta t \epsilon \hat{k}^q + \Delta t \gamma i^{2k} \hat{k}^{2k}). \quad (30)$$

Even in the absence of the spurious growth term $u \Delta t \epsilon \hat{k}^q$, we have $|G| > 1$ due to the first parenthetical term, implying that forward Euler is always unstable on the semi-discrete advection equation. This is because its stability region does not include the imaginary axis. This motivates the use of other Runge-Kutta (RK) methods that contain the imaginary axis. In general, $G(z)$ is a polynomial for all explicit RK methods. For example, for forward Euler, we have

$$G(z) = 1 + z, \quad (31)$$

while for the classical fourth-order RK method (RK4) we have

$$G(z) = 1 + z + \frac{1}{2}z^2 + \frac{1}{6}z^3 + \frac{1}{24}z^4, \quad (32)$$

where $z = -u \Delta t (i\hat{k} - \epsilon \hat{k}^q) + \Delta t \gamma i^{2k} \hat{k}^{2k}$. Requiring $|G(z)| \leq 1$ allows us to compute γ for all RK methods. We demonstrate the calculation of γ for RK4. First, partition $z = -u \Delta t (i\hat{k} - \epsilon \hat{k}^q) + \Delta t \gamma i^{2k} \hat{k}^{2k}$ as $z = z_1 + z_2$, where

$$z_1 = -u \Delta t i \hat{k}, z_2 = u \Delta t \epsilon \hat{k}^q + \Delta t \gamma i^{2k} \hat{k}^{2k}. \quad (33)$$

In this form, it is clear that z_1 arises naturally from the advection equation, while z_2 arises from both the spurious growth mode and the hyperviscosity term. Substituting $z = z_1 + z_2$ into the expression for $G(z)$ for RK4, we obtain

$$G(z) = 1 + z_1 + z_2 + \frac{1}{2}(z_1 + z_2)^2 + \frac{1}{6}(z_1 + z_2)^3 + \frac{1}{24}(z_1 + z_2)^4. \quad (34)$$

Now, partition $G(z)$ as $G(z) = G_1(z_1) + G_2(z_1, z_2)$, where

$$G_1(z_1) = 1 + z_1 + \frac{1}{2}z_1^2 + \frac{1}{6}z_1^3 + \frac{1}{24}z_1^4, \quad (35)$$

$$G_2(z_1, z_2) = z_2 \left[\frac{z_1^3}{6} + \frac{z_1^2}{2} + z_1 + 1 + z_2 \left(\frac{z_1^2}{4} + \frac{z_1}{2} + \frac{1}{2} \right) + z_2^2 \left(\frac{z_1}{6} + \frac{1}{6} \right) + \frac{z_2^3}{24} \right]. \quad (36)$$

In the absence of growth and hyperviscosity terms, stability would require $|G_1(z)| \leq 1$. If $|G(z)| \leq 1$ and $|G_1(z_1)| \leq 1$, it follows that we require $G_2(z_1, z_2) = 0$. Since G_2 may not even be real, the easiest way to compute γ in practice is to enforce $G_2(z_1, z_2) = 0$. In other words, we require that

$$z_2 \left[\frac{z_1^3}{6} + \frac{z_1^2}{2} + z_1 + 1 + z_2 \left(\frac{z_1^2}{4} + \frac{z_1}{2} + \frac{1}{2} \right) + z_2^2 \left(\frac{z_1}{6} + \frac{1}{6} \right) + \frac{z_2^3}{24} \right] = 0. \quad (37)$$

To satisfy this, we set $z_2 = 0$, *i.e.*,

$$u \Delta t \varepsilon \hat{k}^q + \Delta t \gamma i^{2k} \hat{k}^{2k} = 0. \quad (38)$$

Solving (38) for γ , we get

$$\gamma = (-1)^{1-k} \hat{k}^{q-2k} u \varepsilon, \quad (39)$$

The same expression for γ is obtained in the case of all explicit RK methods (not shown). The factor ε , which is the spurious growth mode, can be estimated numerically. We also know that the maximum wavenumber $\hat{k}$ that can be resolved on a regular grid of cell width h is $2h^{-1}$ (for any spatial dimension d). Substituting $2h^{-1}$ in place of $\hat{k}$ gives us

$$\gamma = (-1)^{1-k} 2^{q-2k} h^{2k-q} u \varepsilon. \quad (40)$$

Note that the above expression for γ has an explicit dependence on the velocity u , the largest real part ε , and the growth exponent q . Further, the formula contains a factor of 2^{q-2k} , much like the 2^{-6} factor that appears in [4, 6]. We defer a discussion on q to a later section.

4.1.2. Time-stepping with explicit multistep methods—We next calculate γ for the case of explicit multistep methods. To illustrate the procedure, consider the Adams-Bashforth formula of order 2 (AB2) [35] applied to the generic ODE

$$\frac{\partial c}{\partial t} = \mu c. \quad (41)$$

The AB2 discretization of this ODE is

$$c^{n+1} = c^n + \frac{3}{2} \Delta t \mu c^n - \frac{1}{2} \Delta t \mu c^{n-1}. \quad (42)$$

Now, let $c^{n-1} = \hat{c}(t_{n-1})e^{i\hat{k}x}$ so that $c^n = Gc^{n-1}$ and $c^{n+1} = G^2c^{n-1}$, where G is the growth factor as before. Substituting these equations into the AB2 scheme, we obtain the following quadratic equation for the growth factor G :

$$G^2 - G\left(1 + \frac{3}{2} \Delta t \mu\right) + \frac{1}{2} \Delta t \mu = 0. \quad (43)$$

If AB2 is applied to (27), it follows from (43) that

$$G^2 - G\left(1 + \frac{3}{2}z_1 + \frac{3}{2}z_2\right) + \frac{1}{2}z_1 + \frac{1}{2}z_2 = 0, \quad (44)$$

where $z_1 = -u \Delta t i \hat{k}$ and $z_2 = u \Delta t \epsilon \hat{k}^q + \Delta t \gamma i^{2k} \hat{k}^{2k}$. This equation can be rewritten as

$$\underbrace{G^2 - G\left(1 + \frac{3}{2}z_1\right) + \frac{1}{2}z_1}_{t_1} + \underbrace{z_2\left(\frac{1}{2} - \frac{3}{2}G\right)}_{t_2} = 0. \quad (45)$$

The term t_1 is free of spurious growth (and hyperviscosity), and is the term obtained by applying the AB2 method to the advection equation. If $t_2 = 0$ and a stability criterion on the term $u \hat{k}$ is met, we have $|G| \leq 1$. A straightforward way to obtain an expression for γ is therefore to require t_2 to vanish entirely. More specifically, we require

$$z_2\left(\frac{1}{2} - \frac{3}{2}G\right) = 0. \quad (46)$$

Since G may take on a wide range of values despite the constraint $|G| \leq 1$, the only way to guarantee the above equality is to require $z_2 = 0$. This gives us

$$u \Delta t \epsilon \hat{k}^q + \Delta t \gamma i^{2k} \hat{k}^{2k} = 0, \quad (47)$$

which yields

$$\gamma = (-1)^{1-k} 2^q - 2^k h^{2k} - q u \varepsilon, \quad (48)$$

which is identical to the formula for γ in the RK4 case. This analysis is easily done for all explicit multistep methods, and even for the explicit parts of IMEX multistep methods. For an example of the latter, consider the explicit/extrapolated BDF4 method (the explicit part of the IMEX-BDF4/SBDF4 method [35]) applied to the same ODE:

$$\frac{\frac{25}{12}c^{n+1} - 4c^n + 3c^{n-1} - \frac{4}{3}c^{n-2} + \frac{1}{4}c^{n-3}}{\Delta t} = \mu(4c^n - 6c^{n-1} + 4c^{n-2} - c^{n-3}). \quad (49)$$

The corresponding equation for $G(z)$ is now quartic:

$$G^4 - \frac{48}{25}G^3(1+z) + \frac{36}{25}G^2(1+2z) - \frac{16}{25}G(1+3z) + \frac{3}{25}(1+4z) = 0. \quad (50)$$

Writing $z = z_1 + z_2$ as before and splitting the equation into two terms $t_1(z_1)$ and $t_2(z_2, G)$ yields

$$t_1 = G^4 - \frac{48}{25}G^3(1+z_1) + \frac{36}{25}G^2(1+2z_1) - \frac{16}{25}G(1+3z_1) + \frac{3}{25}(1+4z_1), \quad (51)$$

$$t_2 = z_2 \left(-\frac{48}{25}G^3 + \frac{72}{25}G^2 - \frac{48}{25}G + \frac{12}{25} \right), \quad (52)$$

$$t_1 + t_2 = 0. \quad (53)$$

To have $t_2 = 0$ for stability, we again require $z_2 = 0$. This yields the same expression for γ as in the AB2 and RK4 cases. As an aside, note that it is possible to obtain either $G_2 = 0$ in the RK case or $t_2 = 0$ in the multistep case if the term z_1 or the growth factor G satisfies certain constraints (in addition to the CFL constraint). However, these cases are very rare and obtained only when t , h , and ε take on specific values.

4.2. Selecting γ in the diffusion-dominated regime ($u \rightarrow 0$)

Another important limit is the diffusive regime where v dominates over u . Current RBF-FD methods are not guaranteed to produce purely negative real eigenvalues when approximating the Laplacian. While a mild spread along the imaginary axis is acceptable for parabolic problems, eigenvalues with positive real parts are undesirable as they would cause spurious

solution growth. To calculate the amount of hyperviscosity-based stabilization required to counter this growth, we use the same approach as in Section 4.1, and define the auxiliary 1D PDE

$$\frac{\partial c}{\partial t} = \nu \frac{\tilde{\partial}^2 c}{\partial x^2} + \gamma \frac{\partial^{2k} c}{\partial x^{2k}}, k \in \mathbb{N}, \quad (54)$$

where the operator $\frac{\tilde{\partial}^2}{\partial x^2}$ acting on a plane wave produces the correct decay mode as well as a spurious growth mode, so that if $c(x, t) = \hat{c}(t)e^{ikx}$, $\frac{\tilde{\partial}^2 c}{\partial x^2} = (-\hat{k}^2 + \eta \hat{k}^q)\hat{c}(t)e^{ikx}$, where $\eta \neq 0$ is responsible for the spurious growth. We now discretize (54) in time using the second-order backward difference formula (BDF2) (which is A-stable) [35], obtaining

$$\frac{3c^{n+1} - 4c^n + c^{n-1}}{2 \Delta t} = \nu \frac{\tilde{\partial}^2 c^{n+1}}{\partial x^2} + \gamma \frac{\partial^{2k} c^{n+1}}{\partial x^{2k}}, \quad (55)$$

for which the growth factor G satisfies

$$\left(1 - \frac{2}{3} \Delta t \mu\right) G^2 - \frac{4}{3} G + \frac{1}{3} = 0, \quad (56)$$

where $\mu = \nu(-\hat{k}^2 + \eta \hat{k}^q) + \gamma i^{2k} \hat{k}^{2k}$. Collecting all diffusion terms into term t_1 , and spurious growth and hyperviscosity terms in term t_2 gives:

$$t_1 = \left(1 + \frac{2}{3} \Delta t \hat{k}^2\right) G^2 - \frac{4}{3} G + \frac{1}{3}, t_2 = -\frac{2}{3} \Delta t \left(\nu \eta \hat{k}^q + \gamma i^{2k} \hat{k}^{2k}\right), \quad (57)$$

$$t_1 + t_2 = 0. \quad (58)$$

For stability, we require that $t_2 = 0$. This gives:

$$\nu \eta \hat{k}^q + \gamma i^{2k} \hat{k}^{2k} = 0, \quad (59)$$

$$\Rightarrow \gamma = (-1)^{1-k} 2^{q-2k} h^{2k-q} \nu \eta. \quad (60)$$

If we had lumped the diffusion term into term t_2 , we would have obtained an expression for γ that accounts for the stabilization afforded by natural diffusion. We do not take this approach in this article.

4.3. Selecting γ for coexisting advection and diffusion ($u \neq 0, \nu \neq 0$)

Having examined both the advective and diffusive limits for reasonable choices of time integrators, we now consider the important case where both advection and diffusion may contribute significantly to the transport process. In this scenario, it is common to treat the diffusion term implicitly in time, and the advection term explicitly. Consequently, it is reasonable (in terms of computational efficiency) to also treat the hyperviscosity term implicitly in time. The relevant auxiliary PDE with hyperviscosity is now

$$\frac{\partial c}{\partial t} + u \frac{\partial c}{\partial x} = \nu \frac{\partial^2 c}{\partial x^2} + \gamma \frac{\partial^{2k} c}{\partial x^{2k}}. \quad (61)$$

For plane waves, this yields the ODE

$$\frac{\partial c}{\partial t} = \mu_1 c + \mu_2 c, \quad (62)$$

where

$$\mu_1 = -uik + u\epsilon k^{q_1}, \mu_2 = \nu(-k^2 + \eta k^{q_2}) + \gamma i^{2k} k^{2k}, \quad (63)$$

where q_1 and q_2 are the growth exponents for the approximate gradient and laplacian respectively. To keep our analysis simple, we discretize (62) using the IMEX-BDF2 or SBDF2 method [35]; however, our technique carries over to higher order SBDF methods. The SBDF2 discretization is given by

$$\frac{3c^{n+1} - 4c^n + c^{n-1}}{2\Delta t} = 2\mu_1 c^n - \mu_1 c^{n-1} + \mu_2 c^{n+1}. \quad (64)$$

Substituting in a plane wave expression yields the following equation for the growth factor G :

$$\left(1 - \frac{2}{3}\mu_2 \Delta t\right)G^2 - \frac{4}{3}G(1 + \Delta t\mu_1) + \frac{1}{3}(1 + 2\Delta t\mu_1) = 0. \quad (65)$$

Collecting advection and diffusion terms into the term t_1 , and spurious growth and hyperviscosity terms into the term t_2 , we obtain

$$t_1 = \left(1 + \frac{2}{3}\nu\hat{k}^2 \Delta t\right)G^2 - \frac{4}{3}G(1 - \Delta t u i \hat{k}) + \frac{1}{3}(1 - 2 \Delta t u i \hat{k}), \quad (66)$$

$$t_2 = -\frac{2}{3} \Delta t G^2 (\nu \eta \hat{k}^{q_2} + \gamma i^{2k} \hat{k}^{2k}) + \frac{2}{3} \Delta t u \varepsilon \hat{k}^{q_1} (1 - 2G), \quad (67)$$

$$t_1 + t_2 = 0. \quad (68)$$

As before, we require that $t_2 = 0$, *i.e.*,

$$-G^2 (\nu \eta \hat{k}^{q_2} + \gamma i^{2k} \hat{k}^{2k}) + u \varepsilon \hat{k}^{q_1} (1 - 2G) = 0. \quad (69)$$

Now, let $\gamma = \gamma_1 + \gamma_2$. Then, we have

$$-G^2 (\nu \eta \hat{k}^{q_2} + \gamma_1 i^{2k} \hat{k}^{2k} + \gamma_2 i^{2k} \hat{k}^{2k}) + u \varepsilon \hat{k}^{q_1} (1 - 2G) = 0. \quad (70)$$

We can select γ_2 to cancel out the growth term $\eta \hat{k}^{q_2}$, giving us

$$\gamma_2 = (-1)^{1-k} 2^{q_2-2k} h^{2k-q_2} \eta \nu, \quad (71)$$

and leaving us with

$$-G^2 \gamma_1 i^{2k} \hat{k}^{2k} + u \varepsilon \hat{k}^{q_1} (1 - 2G) = 0. \quad (72)$$

Dividing the above equation by the non-zero quantity $-\gamma_1 i^{2k} \hat{k}^{2k}$, we get the quadratic equation

$$G^2 + \frac{u \varepsilon \hat{k}^{q_1}}{\gamma_1 i^{2k} \hat{k}^{2k}} (2G - 1) = 0. \quad (73)$$

The roots of this equation are given by

$$G = -\frac{u\hat{\epsilon}\hat{k}^{q_1}}{\gamma_1 i^{2k}\hat{k}^{2k}} \pm \frac{1}{\gamma_1 i^{2k}\hat{k}^{2k}} \sqrt{\left(u\hat{\epsilon}\hat{k}^{q_1}\right)\left(u\hat{\epsilon}\hat{k}^{q_1} + \gamma_1 i^{2k}\hat{k}^{2k}\right)}. \quad (74)$$

Stability requires that $|G| \leq 1$. This can be achieved if

$$u\hat{\epsilon}\hat{k}^{q_1} + \gamma_1 i^{2k}\hat{k}^{2k} = 0, \quad (75)$$

which gives

$$\gamma_1 = (-1)^{1-k} 2^{q_1-2k} h^{2k-q_1} u \hat{\epsilon}. \quad (76)$$

Finally, since $\gamma = \gamma_1 + \gamma_2$, we have

$$\gamma = (-1)^{1-k} 2^{q_1-2k} h^{2k-q_1} u \hat{\epsilon} + (-1)^{1-k} 2^{q_2-2k} h^{2k-q_2} \eta \nu. \quad (77)$$

From the analysis, it is clear that the hyperviscosity for spurious growth in advection and diffusion is *additive* under the assumptions of our model. This expression naturally encodes the other cases seen so far. For instance, in the limit $\nu \rightarrow 0$, we recover (40), while in the limit $u \rightarrow 0$, we recover (60). The expressions for γ , though derived using a 1D model, generalize straightforwardly to higher dimensions by replacing u with $\|\mathbf{u}\|_\infty$ (the maximum pointwise spatial velocity or an estimate thereof).

In this article, our simulation results are expressed in terms of the *Peclet number*, which is given $\text{Pe} = \frac{uL^*}{\nu}$, where L^* is some characteristic length scale. To get a higher Peclet number, one can either fix ν and make u larger, or fix u and make ν smaller. The former case is interesting as it magnifies the effect of the spurious eigenvalue ϵ , while the latter serves to reduce any stabilizing diffusion. In Section 6, we focus on the case of increasing u for fixed ν , as this is a more effective test of our hyperviscosity formulation. However, we have observed that our hyperviscosity formulation is extremely robust and stable in the scenario of decreasing ν as well; this can be seen in Section 7. In both these sections, we demonstrate that it is perfectly safe to use (40) for γ in the context of IMEX multistep methods.

4.4. Estimating growth exponents (q_1 and q_2)

Thus far, our growth models have only included a *generic* spurious growth term of the form $\epsilon \hat{k}_1^q$ or $\eta \hat{k}_2^q$. To the best of our knowledge, the literature lacks any theoretical estimates for q_1 or q_2 on any given node set. Our goal in this section is to give a simple and cost-effective *numerical* procedure for computing these values on a given node set. Without loss of

generality, we proceed with a description of this procedure on a 1D node set. Let $X = \{x_j\}_{j=1}^N$ be a set of nodes on which we obtain the discrete RBF-FD differentiation matrix G^x that approximates $\frac{\partial}{\partial x}$ in the advection term. Our goal now is to estimate q_1 for the matrix G^x ; the procedure for q_2 is identical. Consider the function $f(x) = e^{\hat{k}x}$. Its derivative is given exactly by $g(x) = \hat{k}e^{\hat{k}x}$. Then, define

$$\underline{f} = f(x)|_X, \underline{g} = g(x)|_X = \hat{k}\underline{f}, \quad (78)$$

where $\underline{f}, \underline{g}$ are vectors with N entries. On the other hand, the *approximate* derivative of $f(x)$ on the node set X is given by

$$\underline{\tilde{g}} = G^x \underline{f}. \quad (79)$$

According to our growth model, we know that G^x is represented by the auxiliary differential operator $\frac{\tilde{\partial}}{\partial x}$. Consequently, ignoring truncation errors in G^x , the growth model gives

$$\underline{\tilde{g}} = \left(\hat{k} - \varepsilon \hat{k}^{q_1} \right) \underline{f}. \quad (80)$$

This allows us to write

$$\| \underline{g} - \underline{\tilde{g}} \| = \varepsilon \hat{k}^{q_1} \| \underline{f} \|. \quad (81)$$

In the scenario with spurious growth modes, we know that $\varepsilon \| \underline{f} \| \rightarrow 0$. Dividing by $\varepsilon \| \underline{f} \|$ and reversing the sides of the equation for clarity, we get

$$\hat{k}^{q_1} = \frac{\| \underline{g} - \underline{\tilde{g}} \|}{\varepsilon \| \underline{f} \|}, \quad (82)$$

$$\Rightarrow q_1 = \frac{\ln(\| \underline{g} - \underline{\tilde{g}} \|) - \ln(\varepsilon \| \underline{f} \|)}{\ln \hat{k}}. \quad (83)$$

As we remarked when deriving γ , the largest wavenumber $\hat{k}$ that can be represented on a grid of cell width h is $2h^{-1}$. Thus, in the above formula for q_1 , we use $\hat{k} = 2h^{-1}$, where h is taken to be $h = N^{-1/d}$. To fully extend the above definition of q_1 to dimension $d > 1$, we need only change the definition of $f(x)$ to be a product of exponentials in each spatial variable.

The above procedure is repeated for each of the gradient differentiation matrices G^x, G^y , and G^z , and the largest q_1 value of the three is used. While any function $f(x)$ could have been chosen in principle, the use of plane waves allows a direct comparison to our growth model and an easy evaluation on any node set. In addition, to compute q_2 , we use $g(x) = -\hat{k}^2 e^{i\hat{k}x}$ as our reference derivative function and $\tilde{g} = \left(-\hat{k}^2 + \eta \hat{k}^{q_2}\right) \underline{f}$, but the above formula remains the same (with η in place of ε). In practice though, since $\eta \approx 0$ and we use implicit time-stepping, it is typically unnecessary to compute q_2 .

To complete the above discussion, we note that ε and η must be estimated to find q_1 and q_2 . These factors correspond naturally to the real part of the eigenvalue with the largest real part. If a loose tolerance is used in the eigenvalue computation (as is done in this article), the estimates for q_1 and q_2 are likely to be an approximation as well. Note also that our definition allows q_1 and q_2 to be negative, though in practice we have observed that q_1 and q_2 are typically between 0 and 1. The approach detailed in this section *eliminates* all free parameters from the expression for γ for the cost of d sparse matrix-vector multiplications for estimating q , and a few sparse matrix-vector multiplications for estimating ε . In contrast, expressions for γ in the literature contain factors that must be tuned to achieve stability [4, 10]. If theoretical developments should ever give *a priori* estimates for ε and η , these can be substituted directly into the formula for γ .

4.5. Spurious modes in the hyperviscosity operator

Our analysis, for convenience, has neglected the fact that an RBF-FD approximation to the operator $\frac{\partial^{2k}}{\partial x^{2k}}$ may itself contain spurious eigenvalues that cause undesirable growth or decay (so that they always produce spurious growth when we consider γ^{-k} since $\text{sign}(\gamma) = (-1)^{1-k}$). It is reasonable to wonder how the expression for γ would change if one factored in this additional spurious mode. Fortunately, our analysis technique from the previous sections is perfectly amenable to this case as well. Let $\frac{\tilde{\partial}^{2k}}{\partial x^{2k}}$ be the *auxiliary* hyperviscosity operator, defined by its action on plane waves:

$$\frac{\tilde{\partial}^{2k} e^{i\hat{k}x}}{\partial x^{2k}} = \left(i^{2k} \hat{k}^{2k} \pm \tau \hat{k}^{q_3}\right) e^{i\hat{k}x}, \quad (84)$$

where $\tau \neq 0$, but the spurious term has the opposite sign of the eigenvalues of the operator $\frac{\partial^{2k}}{\partial x^{2k}}$. Without loss of generality, we rederive γ for explicit RK4 for this new scenario. Using our previous notation of $G(z)$ to represent the growth factor, splitting $z = z_1 + z_2$, where z_1 contains the advection term, and z_2 contains *all* spurious growth terms, we again require $z_2 = 0$ to obtain stability. This yields the condition

$$u \Delta t \varepsilon \hat{k}^{q_1} + \Delta t \gamma \left(i^{2k} \hat{k}^{2k} \pm \tau \hat{k}^{q_3}\right) = 0, \quad (85)$$

where the symbols have their usual meanings. Re-arranging, we get

$$\gamma = - \frac{u \varepsilon \hat{k}^{q_1}}{i^{2k} \hat{k}^{2k} \pm \tau \hat{k}^{q_3}}, \quad (86)$$

$$\Rightarrow \gamma = \frac{(-1)^{1-k} u \varepsilon 2^{q_1-2k} \frac{h^{2k-q_1}}{h}}{1 \pm (-1)^k \tau 2^{q_3-2k} \frac{h^{2k-q_3}}{h}}, \quad (87)$$

for $\hat{k} = 2h^{-1}$. Since the τ term is scaled by the typically very small value of $2^{q_3-2k} h^{2k-q_3}$, we find for the experiments in this article that it is safe to neglect the effect of any spurious modes in the RBF-FD approximation to the $\hat{k}$ operator when computing γ .

4.6. Selecting k

We next turn to the selection k , the power of the Laplacian in the term γ^{-k} . It has already been observed in the RBF-FD literature that k must be increased as the stencil size n is increased [10]. However, the formula presented in [10] was heuristic and was fine-tuned to the node sets on the sphere. In more modern RBF-FD formulas on Euclidean domains [4], this approach appears to have been tested and discarded as unnecessary. However, in our experiments, we found that it was indeed necessary to scale k with n on scattered nodes in domains with irregular boundaries.

We adopt an approach based on the spectral methods literature. Specifically, examining Ma's work on the Chebyshev-Legendre super viscosity method [25, 27], we see in Remark 2 below (5.21) that the order k (s in their notation) of the spectral superviscosity must scale as

$$k \leq O(\ln N), \quad (88)$$

where N is the total number of points. Our recipe for k uses a modified version of this scaling law. Bearing in mind that an FD method reproduces a spectral method as the stencil size $n \rightarrow N$, this prompts us to select

$$k = \lfloor 1.5 \ln n \rfloor. \quad (89)$$

This formula has the effect of increasing k with n (as desired). The actual constant in front of the $\ln n$ term does not appear to have a great impact on accuracy. While the logarithmic scaling may arise naturally as a consequence of the non-uniform resolution of Chebyshev-Legendre spectral methods, it has benefits when used with RBF-FD approximations also. First, the law (89) is easy to compute and requires no fine tuning. Second, it serves to increase k slowly with n . This slow increase ensures that the stencil sizes and polynomial degrees needed for approximating $\hat{k}$ also grow only slowly with n . Together, the formulas

for γ and k give us a hyperviscosity operator that produces the correct damping behavior in the $n \rightarrow N$ limit, and vanishes in the $h \rightarrow 0$ limit, all without any free parameters.

4.7. RBF approximations to the hyperviscosity operator

As in the case of the other terms in (23), we approximate the hyperviscosity term with PHS RBFs augmented with polynomials. First, the PHS RBF used for hyperviscosity must have sufficient smoothness so as to be able to apply the operator γ^{-k} to it. We use the PHS RBF

$$\phi_{hyp}(r) = r^{2k+1}, \quad (90)$$

which is the PHS RBF of minimal smoothness required. It follows that, $\gamma^k \phi_{hyp}(r) = Cr$, where C is a known constant involving k and the spatial dimension d . This follows from repeated application of the formula for the Laplacian of a radial function in d -dimensions: $r^m = m(m-d+2)r^{m-2}$.

In addition to selecting the PHS RBF, we must also select the polynomial degree ℓ_{hyp} for approximating the hyperviscosity operator. Our approach here is very simple. For a PHS RBF of degree $2k+1$, the minimal polynomial degree required to prove unisolvency of the RBF interpolant is k [16]. We thus pick $\ell_{hyp} = k$. The observant reader may be concerned: after all, for even a first-order approximation of the operator γ^{-k} , it would seem that $\ell_{hyp} = 2k$ is the correct choice. However, this ignores the fact that the operator in question is γ^{-k} . Since $\gamma = O(h^{2k-q})$ where we typically have $q = O(1)$, the scaling with γ makes our approximation to this operator very high-order. When computing overlapped RBF-FD weights for the operator γ^{-k} , the choice of $\ell_{hyp} = k$ also zeros out the derivatives of polynomial terms upto degree ℓ which in our case is *all* terms.

The different parameters and basis function choices for our RBF-FD methods are summarized in Table 1.

5. Ghost node formulation

5.1. Method of lines

In this section, we discuss our technique for time integration of (23). In addition to selecting a time integrator, this requires the use of ghost nodes in the spatial discretization of the domain Ω . As mentioned in Section 1, we observed that ghost nodes were required to enhance the global stability of our RBF-FD discretizations. This is not without precedent in the RBF-FD literature [4–6], though it is still an open question whether they are required for all problems [23]. We will now focus on solving (23) with the inclusion of a forcing term and boundary conditions:

$$\frac{\partial c}{\partial t} = f(c) - \mathbf{u} \cdot \nabla c + \nu \Delta c + \gamma \Delta^k c, \mathbf{x} \in \Omega, \quad (91)$$

$$\mathcal{B}c = g(\mathbf{x}), \mathbf{x} \in \partial\Omega, \quad (92)$$

where f is some forcing or reaction term. Our approach is to discretize the above equations in a method of lines formulation: we first discretize space, then discretize time.

We assume now that the node set X contains the required ghost nodes; see [33, 36] for some ways to generate node sets and ghost nodes. Partition X as $X = X_i \cup X_b \cup X_g$, where the subscripts i, b , and g indicate interior, boundary and ghost nodes, respectively. Let C be the numerical solution to (23); write the vector C as $[C_i, C_b, C_g]^T$. Let L be the discrete Laplacian, G^x , G^y , and G^z the discrete components of the gradient operator, and H be the discrete hyperviscosity operator, all formed using overlapped RBF-FD on the node set. Let $\mathbf{u} = [u_x, u_y, u_z]^T$. For each discrete differential operator K , we require the following matrices:

$$K_{if} = [K_{ii} \quad K_{ib} \quad K_{ig}], \quad (93)$$

$$K_{bf} = [K_{bi} \quad K_{bb} \quad K_{bg}], \quad (94)$$

where K_{ii} maps vectors from the interior to the interior, K_{ib} from the boundary to the interior, and so forth. Finally, let B be the discrete boundary condition operator approximating $\mathcal{B}$. We can now write the semi-discrete analog of (92) as

$$\frac{\partial C_i}{\partial t} = f(C_i) - (u_x)_i^T G_{if}^x C - (u_y)_i^T G_{if}^y C - (u_z)_i^T G_{if}^z C + \nu L_{if} C + \gamma H_{if} C, \quad (95)$$

$$\frac{\partial C_b}{\partial t} = f(C_b) - (u_x)_b^T G_{bf}^x C - (u_y)_b^T G_{bf}^y C - (u_z)_b^T G_{bf}^z C + \nu L_{bf} C + \gamma H_{bf} C, \quad (96)$$

$$B_{bi} C_i + B_{bb} C_b + B_{bg} C_g = g_b, \quad (97)$$

where $g_b = g(\mathbf{x})|_{X_b}$. The above system enforces the PDE up to and including the boundary, and enforces boundary conditions at the boundary to accommodate the additional unknowns at the ghost nodes (of which there are as many as boundary nodes). Next, we proceed with the time-discretization of the constrained ODE system (95)–(97). To simplify the discussion, we use forward Euler for all nonlinear terms, and backward Euler for stiff terms (including hyperviscosity). The discrete system of equations is:

$$\begin{aligned} \frac{C_i^{n+1} - C_i^n}{\Delta t} &= f(C_i^n) - \left((u_x)_i^n \right)^n G_{if}^x C^n - \left((u_y)_i^n \right)^n G_{if}^y C^n - \left((u_z)_i^n \right)^n G_{if}^z C^n + \nu L_{if} C^{n+1} \quad (98) \\ &+ \gamma H_{if} C^{n+1}, \end{aligned}$$

$$\begin{aligned} \frac{C_b^{n+1} - C_b^n}{\Delta t} &= f(C_b^n) - \left((u_x)_b^n \right)^n G_{bf}^x C^n - \left((u_y)_b^n \right)^n G_{bf}^y C^n - \left((u_z)_b^n \right)^n G_{bf}^z C^n + \nu L_{bf} C^{n+1} \quad (99) \\ &+ \gamma H_{bf} C^{n+1}, \end{aligned}$$

$$B_{bf} C_i^{n+1} + B_{bb} C_b^{n+1} + B_{bg} C_g^{n+1} = g_b^{n+1}. \quad (100)$$

Higher order analogues follow naturally. We use the semi-implicit backward difference formula of order 4 (SBDF4), started with a step each of SBDF1, SBDF2, and SBDF3 [35]. Examining the above systems of equations, one can see that the ghost point values C_g^{n+1} are a part of the system of unknowns, and that the ghost point values appear on the right hand side at time level n in the first-order scheme, and time levels n and $n-1$ in the second-order scheme. Clearly, except for the first time-step, there is no need to fill the ghost cells. However, the first time-step requires some careful attention in the case of a general boundary operator $\mathcal{B}$.

If $\mathcal{B}$ enforces Neumann or Robin conditions, the ghost values C_g^n can be obtained by solving the following linear system [4]:

$$B_{bi} C_i^n + B_{bb} C_b^n + B_{bg} C_g^n = g_b^n, \quad (101)$$

$$\Rightarrow B_{bg} C_g^n = g_b^n - (B_{bi} C_i^n + B_{bb} C_b^n). \quad (102)$$

However, in the case of pure Dirichlet boundary conditions (as opposed to mixed, Neumann, or Robin conditions), the matrix B_{bg} is a matrix of zeros and therefore has no inverse. A simple solution that works for a variety of boundary condition types is to spatially extrapolate the solution from the interior and boundary to the ghost points using local RBF interpolants of nearby nodal values. This is only done once (at time $t=0$), since all subsequent ghost node values are obtained from the time-stepping process itself.

5.2. Positivity-preserving filter

Many of our target mathematical models require that the values of the transported variable c be positive at all times, with even small spurious negative values potentially leading to enormous spurious feedback loops [37, 38]. High-order RBF-FD methods do exhibit a small amount of dispersion; see the figures in [4]. While a full exploration of filters to rectify this dispersion is beyond the scope of this work, we adopt the following simple approach to ensure positivity:

1. Check if the elements of the solution vectors C_j and C_b have a negative sign.
2. Set any element with a negative sign to zero.

This filter is applied at the end of each time-step. As we will show in the results section, this does not affect convergence rates at all in the context of advection-diffusion problems. Further, it appears to reduce errors and restore correct convergence rates on a pure advection test case. For more details, see Section 6.

6. Results

We now test our hyperviscosity formulation via convergence studies on the forced advection-diffusion equation over a range of Peclet numbers. The forcing term is selected to maintain a prescribed solution for all time, and the prescribed solution is used to test spatial convergence rates. We solve this test problem on the closed unit disk in $\mathbb{R}^2$ and the closed unit ball in $\mathbb{R}^3$. Node sets were generated using the $O(N)$ node generation algorithms described in [33]. Some examples of these node sets are shown in Figure 4. Given a true solution $c(\mathbf{x}, t)$ and a numerical solution $C(\mathbf{x}, t)$, we compute relative errors at the final time $t = 2$ on the node set X as $e_{\ell_p} = \frac{\|c_X - C_X\|_p}{\|c_X\|_p}$, where $p = 2, \infty$.

6.1. Advection on the unit disk

Our first test involves applying our new hyperviscosity formulation to solve a pure advection problem on the unit disk. Though the focus of this article is on advection-diffusion equations, this test problem allows us to study stability and convergence rates without any stabilizing diffusion in the problem. Further, it allows us to test the hyperviscosity formulation from Section 4.1.1 in the context of explicit RK methods. Finally, we also use this test to determine the effect of filtering on convergence rates.

The test involves advecting a Gaussian in a rotational velocity field. The initial condition is given by

$$c(x, y, 0) = e^{-\frac{(x-x_0)^2 + (y-y_0)^2}{\lambda^2}}, \quad (103)$$

where $x_0 = 0.5$, $y_0 = 0.05$, and $\lambda = 1/8$. The velocity field is given by

$$\mathbf{u}(x, y, t) = 2\pi \sin(\pi t)[y, -x]. \quad (104)$$

This velocity field effectively moves the initial condition some distance, and then brings it back to its starting position at time $t = 2$, which is our chosen final time. Since the problem is a fully explicit advection equation of the form $\frac{\partial c}{\partial t} = -\mathbf{u} \cdot \nabla c$, an excellent time integrator is the explicit (classical) RK4 method. In this case, since we only have a first order differential operator in the PDE ($\theta = 1$), we set $\xi = \ell$. The results with and without filtering are shown in Figure 5. From Figure 5a, we can see that increasing ℓ increases the convergence rate, except in the case of $\ell = 5$, where we see a reduction in order of convergence. In Figure 5b, we see clearly that using a filter to preserve positivity has no effect on the convergence rates when $\ell = 3, 4$. However, for $\ell = 5$, we see that using the filter restores the correct convergence rate. Our results also demonstrate that our hyperviscosity formulation results in stable simulations in the context of RK4 timestepping, even when the hyperviscosity is stepped explicitly in time (and without any stabilizing diffusion). We use the filter in all our test results, though the errors without the filters appear to be very similar when some diffusion is present.

6.2. Forced Advection-Diffusion on the disk

The next test involves solving the advection-diffusion equation on the unit disk. The manufactured solution is given by

$$c(\mathbf{x}, t) = 1 + \sin(\pi x)\cos(\pi y)e^{-\pi t}, \quad (105)$$

and the incompressible velocity field $\mathbf{u}(\mathbf{x}, t) = [u, v]$ is given by

$$\mathbf{u}(\mathbf{x}, t) = \sin(\pi \|\mathbf{x}\|_2^2)\sin(\pi t)[y, -x]. \quad (106)$$

This velocity field vanishes on the boundary of the disk, allowing us to safely impose Neumann boundary conditions on c there using the manufactured solution. This setup allows us to measure errors in our numerical solution against the prescribed c . To test our hyperviscosity formulation across a range of Peclet numbers, we scale the velocity field by the Peclet number Pe . In all tests shown here, we set the diffusion coefficient to $\nu = 1$. We set the time-step to

$$\Delta t = \frac{h}{2 \|\mathbf{u}\|_{\max}}, \quad (107)$$

where $\|\mathbf{u}\|_{\max}$ is calculated by taking the maximum of $\|\mathbf{u}\|$ over space and time. We simulate the PDE to time $t = 2$ using the SBDF4 method, starting up with a single step each of SBDF1, SBDF2, and SBDF3. The PDE contains both a first-order ($\theta = 1$) and second-order ($\theta = 2$) differential operator; rather than using different polynomial degrees ℓ for each

operator, we use $\theta = 2$ to select $\mathcal{L}$ for a given value of ξ (the desired convergence rate). The results for $\ell = 4, 5, 7$ (with $\xi = \ell - 1$) are shown in Figure 6, plotted as a function of $\sqrt{N}$ (proportional to $1/h$).

First, it is clear from Figure 6 that increasing ξ and therefore $\mathcal{L}$ indeed has the effect of increasing convergence rates in general regardless of Pe . Figure 6a shows an anomalous result for $\ell = 7$ where the convergence rate drops off; however, we have verified that this is due to temporal errors dominating the total error. Indeed, at this spatial resolution and value of Pe , the time-step is roughly $\Delta t = O(10^{-2})$, which incurs a fairly large error even with a fourth-order time-integrator. As Pe is increased and smaller Δt values are consequently used, we see high order convergence rates being restored. However, as the Peclet number is increased, we see an increase in the observed order of convergence. This is easily explained: as Pe goes up, the influence of the diffusion term (with $\theta = 2$) goes down. Consequently, most of the spatial errors come from the gradient operator ($\theta = 1$) since the amount of added hyperviscosity is very small. The polynomial degree $\mathcal{L}$ however, was selected with $\theta = 2$ in mind; it is thus reasonable to expect an extra order of convergence when $\theta = 1$. We note that similar results were seen when the magnitude of the velocity was fixed and ν was decreased to increase Pe (not shown); in those cases, as ν was decreased, it was important to scale the boundary rows of the time-stepping matrix (and the corresponding right hand sides) by $(\Delta t)^{-2}$ to ensure that the spectral radius was less than one. It appears that our proposed formula for γ is stable across a range of parameters and Peclet numbers, as is our proposed scaling law for m and ℓ .

6.3. Forced Advection-Diffusion in the ball

Next, we conducted the 3D analogue of the test on the disk: a convergence study on the forced advection-diffusion equation in the unit ball. In this case, our manufactured solution is

$$c(\mathbf{x}, t) = 1 + \sin(\pi x)\cos(\pi y)\sin(\pi z)e^{-\pi t}, \quad (108)$$

and the incompressible velocity field $\mathbf{u}(\mathbf{x}, t) = (u, v, w)$ is given by

$$\mathbf{u}(\mathbf{x}, t) = \sin(\pi \|\mathbf{x}\|_2^2)\sin(\pi t)[-yz, 2xz, -xy]. \quad (109)$$

It is easy to show that $\mathbf{u}(\mathbf{x}, t)$ vanishes on the unit sphere, allowing us to impose Neumann boundary conditions on $c(\mathbf{x}, t)$ there. The time-steps and Peclet numbers were chosen as in the 2D case, and we set $\nu = 1$ for all tests. We simulate the PDE to time $t = 2$, and measure errors against the manufactured solution. The results are shown in Figure 7, plotted as function of $\sqrt[3]{N}$ (proportional to $1/h$). Once again, as the Peclet number is increased, we see an increase in convergence rates, though more erratically than in the 2D case. For $\xi = 3$, for instance, $Pe = 1, 10$ give similar results, while $Pe = 100$ gives a reduced order of convergence and $Pe = 1000$ gives a much higher convergence rate than expected. However, in general, as $\mathcal{L}$ is increased, we see an increase in convergence rates and a decrease in errors.

7. A coupled problem in the spherical shell

We now apply our method to solving a coupled problem in the spherical shell (inspired by platelet aggregation and coagulation). In this problem, we track a chemical concentration $c(\mathbf{x}, t)$ in a fluid inside a spherical shell domain (inner radius 0.3, outer radius 1.0). $c(\mathbf{x}, t)$ is transported by the advection-diffusion equation in an incompressible velocity field, once again giving:

$$\frac{\partial c}{\partial t} = \mathbf{u} \cdot \nabla c = \nu \Delta c + f_1(\mathbf{x}, t), \mathbf{x} \in \Omega, \quad (110)$$

where $f_1(\mathbf{x}, t)$ is some forcing term, and Ω is the region between the two spheres of radius 0.3 and 1.0. On the outer spherical boundary $\mathbb{S}_1$, c satisfies a time-varying inhomogeneous Neumann boundary condition:

$$-\nu \frac{\partial c(\mathbf{x}, t)}{\partial \mathbf{n}} = g(\mathbf{x}, t), \mathbf{x} \in \mathbb{S}_1, \quad (111)$$

which corresponds to an *inward* flux of c . The inner spherical boundary $\mathbb{S}_2$ is viewed as a reactive “zone” on which chemicals can bind, unbind, and participate in other reactions. We track the *bound* chemical surface density separately, and label it C_B . We assume that at each point on the inner sphere, C_B satisfies the following ODE:

$$\frac{\partial C_B}{\partial t} = k_{on}(C^{Tot} - C_B)c_{amb} - k_{off}C_B + k_{self}C_B(C^{Tot} - C_B) + f_2, \quad (112)$$

where k_{on} and k_{off} are the binding and unbinding rates of C_B , k_{self} is the rate at which C_B reacts with itself, C^{Tot} is the total density of binding sites at each point of the reactive zone, f_2 is some forcing term, and c_{amb} is the concentration of the chemical c at the reactive zone, *i.e.*, $c_{amb} = c(\mathbf{x}, t) |_{\mathbb{S}_2}$. Balancing fluxes at the interface $\mathbb{S}_2$ yields the following time-varying Robin boundary condition on $c(\mathbf{x}, t)$:

$$-\nu \frac{\partial c(\mathbf{x}, t)}{\partial \mathbf{n}} + k_{on}(C^{Tot} - C_B)c(\mathbf{x}, t) = k_{off}C_B, \mathbf{x} \in \mathbb{S}_2. \quad (113)$$

7.1. A manufactured solution to the coupled problem

The functions f_1 , f_2 , and g are usually specified by platelet aggregation and coagulation models. In this article, however, we use these terms to test convergence of our numerical methods on the model problem. We again use the method of manufactured solutions, and use the terms f_1 , and f_2 to make the solution hold true. The procedure is as follows. First, we set $c(\mathbf{x}, t)$ to be

$$c(\mathbf{x}, t) = c(x, y, z, t) = 1 + \sin(\pi x)\cos(\pi y)\sin(\pi z)e^{-\pi t}. \quad (114)$$

Next, we manufacture an incompressible velocity field in the spherical shell. The velocity field $\mathbf{u}(\mathbf{x}, t) = (u, v, w)$ is given by

$$\mathbf{u}(\mathbf{x}, t) = \phi(\mathbf{x})\sin(\pi t)[-yz, 2xz, xy], \quad (115)$$

where $\phi(x, y, z)$ is given by

$$\phi(\mathbf{x}) = \sin(\pi \|\mathbf{x}\|_2^2)\sin\left(\frac{\pi}{0.3^2} \|\mathbf{x}\|_2^2\right). \quad (116)$$

It is easily verified that $\mathbf{u}(\mathbf{x}, t)$ is incompressible in the spherical shell, and satisfies no-slip conditions on both $\mathbb{S}_1$ and $\mathbb{S}_2$; in fact, the shell domain was chosen primarily due to the ease of manufacturing this velocity field. Once $c(\mathbf{x}, t)$ and $\mathbf{u}(\mathbf{x}, t)$ are specified, we compute the forcing term $f_1(\mathbf{x}, t)$ as

$$f_1(\mathbf{x}, t) = \frac{\partial c}{\partial t} + \mathbf{u} \cdot \nabla c - \nu \Delta c. \quad (117)$$

The boundary condition function $g(\mathbf{x}, t)$ is then obtained by applying the Neumann operator $-\nu \frac{\partial}{\partial n}$ to $c(\mathbf{x}, t)$. The function C_B which is consistent with the specified c is obtained by solving (113) for C_B to obtain

$$C_B = \frac{-\nu \frac{\partial c}{\partial n} + k_{on} C^{Tot} c}{k_{on} c + k_{off}}. \quad (118)$$

Using this C_B , we compute the forcing term f_2 as

$$f_2 = \frac{\partial C_B}{\partial t} - k_{on}(C^{Tot} - C_B)c_{amb} + k_{off}C_B - k_{self}C_B(C^{Tot} - C_B). \quad (119)$$

We compute errors against the exact C_B and $c(\mathbf{x}, t)$ in a spatial refinement study.

7.2. Operator splitting for the coupled problem

The fluid-phase chemicals $c(\mathbf{x}, t)$ are coupled to the bound chemicals C_B nonlinearly through the c_{amb} term in (112) and the boundary conditions in (113). In order to efficiently simulate the model equations, we present a simple operator splitting algorithm based on Strang-Marchuk splitting [39, 40]. For convenience, in this section we will abuse notation and use

C_B and c to represent the numerical solutions. Given: C_B^n, c^n (solutions at time level n), the algorithm is as follows:

1. Average C_B^n and C_B^{n-1} to obtain $C_B^{n+1/2}$.
2. Use $C_B^{n+1/2}$ and C_B^n to obtain $C_B^{n+1/2}$ by discretizing (112) with the second-order Adams-Bashforth method (AB2).
3. Compute an approximation $\tilde{C}_B^{n+1}$ to C_B^{n+1} by the locally-third-order extrapolation formula $\tilde{C}_B^{n+1} = \frac{8}{3}C_B^{n+1/2} - 2C_B^n + \frac{1}{3}C_B^{n-1}$.
4. Use $\tilde{C}_B^{n+1}$ to obtain the boundary conditions for c^{n+1} on $\mathbb{S}_2$ as:

$$-\nu \frac{\partial c^{n+1}}{\partial n} + k_{on}(C^{Tot} - \tilde{C}_B^{n+1})c^{n+1} = k_{off}\tilde{C}_B^{n+1}. \quad (120)$$

5. Update c^n to c^{n+1} by discretizing (110) with overlapped RBF-FD in space and the SBDF2 scheme in time (adding hyperviscosity implicitly in time).
6. Update $C_B^{n+1/2}$ to C_B^{n+1} by computing using $c_{amb}^{n+1/2} = 0.5(c^{n+1} + c^n)|_{\mathbb{S}_2}$ in (112) and the AB2 discretization.

It is easy to show that the splitting error is $\mathcal{O}(\tau^2)$, and that the overall error in time is also $\mathcal{O}(\tau^2)$. We leave the investigation of higher-order splitting schemes to future work. For the first step, we use forward Euler in place of AB2, SBDF1 in place of SBDF2, and the locally second-order extrapolation $\tilde{C}_B^1 = 2C_B^{1/2} - C_B^0$ in place of the third-order extrapolation. We fill the initial ghost nodes using the RBF-based spatial extrapolation scheme mentioned previously.

The time-varying boundary conditions on the inner boundary $\mathbb{S}_2$ require care for efficient time-stepping. Since the boundary condition *operator* itself is changing as C_B changes, the entire time-stepping matrix cannot be decomposed and used in a sparse direct solve (as in the previous section). We precompute the rows of the time-stepping matrix corresponding to the domain interior Ω and the outer boundary $\mathbb{S}_1$, and *append* and modify the rows corresponding to the domain boundary $\mathbb{S}_2$ every step. Since the location of $\mathbb{S}_2$ is fixed, we precompute the RBF-FD interpolation matrices and only modify the right hand sides of the linear systems corresponding to the boundary RBF-FD weights. We use the GMRES method to solve the sparse linear system for c [41, 42]. GMRES typically requires a good preconditioner for fast convergence. Fortunately, our problem suggests a natural preconditioner. First, we precompute an auxiliary time-stepping matrix corresponding to Neumann boundary conditions on all boundaries. We then form the sparse incomplete LU (ILU) factorization of this matrix, and use it to precondition GMRES on the problem where

Robin boundary conditions are used. This allowed GMRES to converge in 2-6 iterations on average.

7.3. Results

To fully test our hyperviscosity formulation, we now set the diffusion coefficient to $\nu = 0.01$, and scale the velocity field $\mathbf{u}$ to obtain a Peclet number of $Pe = 1000$. We once again select the time-step as in the advection-diffusion test cases. We do not use any near-boundary refinement in our node sets as this may affect the CFL condition for our problem. We run convergence studies using $\xi = 2, 4$ and 6 , and measure the relative errors in both the approximations to c and C_B . The results are shown in Figure 8.

Figure 8a shows that the spatial error in the numerical approximation to c decreases at the rate of approximately $\xi = \ell - 1$ despite the use of a second-order time-stepping scheme with time-steps close to the CFL constraint. Figure 8b shows that the error in the approximation to C_B decreases at similar rates for $\ell = 3, 7$, and a slightly slower rate for $\ell = 5$. At first thought, this may seem unusual since C_B is governed by an ODE. However, due to our manufactured solution, C_B is affected by a spatial forcing term, and thus can be expected to converge at similar rates as c .

As an aside, we also remark that our discretization scheme is very well-suited to coupled problems of this sort, since our schemes automatically solve for boundary values of c thereby giving us c_{amb} .

8. Summary and Future Work

In this article, we presented high-order numerical schemes for simulating the advection-diffusion equation using the overlapped RBF-FD method. We stabilized our methods using artificial hyperviscosity of the form γ^{-k} . Unlike previous attempts in the RBF-FD literature, our expression for γ was derived using a novel 1D Von-Neumann type analysis based on an explicit representation of spurious growth modes through auxiliary differential operators; the spurious growth mode was estimated numerically using a single matrix-vector multiply per differentiation matrix. We also presented an expression for k based on the spectral methods literature. Using a novel ghost node formulation for IMEX time-stepping, we demonstrated high-order convergence rates on both 2D and 3D advection-diffusion problems. In addition, we demonstrated high-order convergence rates on a more complicated coupled problem in the spherical shell.

While our derived expressions for γ are free of any tuning parameters, they require knowledge about the spectra of the differential operators being stabilized. Specifically, the quantities ε and η must be given as an input to the algorithm. Theoretical foundations for the influence of node sets, m , and ℓ on the values of ε and η are non-existent, forcing us to select these quantities numerically using very loose tolerances (for efficiency). We plan to explore such relationships between these parameters in future work. Also, while our article briefly explores the use of filters to preserve positivity in the numerical solution, a detailed exploration of different types of filters will be the subject of future work. Finally, our goal is to apply our hyperviscosity formulation in the setting of PDEs on time-varying domains.

Acknowledgments

This work was supported in part on NSF grants DMS-1160432 and DMS-1521748 and NIH grant 1R01HL120728.

References

- [1]. Bayona V, Moscoso M, Carretero M, Kindelan M, RBF-FD formulas and convergence properties, *J. Comput. Phys.* 229 (22) (2010) 8281–8295.
- [2]. Davydov O, Oanh DT, Adaptive meshless centres and RBF stencils for poisson equation, *J. Comput. Phys.* 230 (2) (2011) 287–304.
- [3]. Wright GB, Fornberg B, Scattered node compact finite difference-type formulas generated from radial basis functions, *J. Comput. Phys.* 212 (1) (2006) 99–123.
- [4]. Flyer N, Barnett GA, Wicker LJ, Enhancing finite differences with radial basis functions: Experiments on the Navier-Stokes equations, *J. Comput. Phys.* 316 (2016) 39–62.
- [5]. Flyer N, Fornberg B, Bayona V, Barnett GA, On the role of polynomials in RBF-FD approximations: I. Interpolation and accuracy, *J. Comput. Phys.* 321 (2016) 21–38.
- [6]. Barnett GA, A robust rbf-fd formulation based on polyharmonic splines and polynomials, Ph.D. thesis, University of Colorado Boulder (2015).
- [7]. Flyer N, Wright GB, Transport schemes on a sphere using radial basis functions, *J. Comput. Phys.* 226 (2007) 1059–1084.
- [8]. Flyer N, Wright GB, A radial basis function method for the shallow water equations on a sphere, *Proc. Roy. Soc. A* 465 (2009) 1949–1976.
- [9]. Fornberg B, Lehto E, Stabilization of RBF-generated finite difference methods for convective PDEs, *J. Comput. Phys.* 230 (2011) 2270–2285.
- [10]. Flyer N, Lehto E, Blaise S, Wright GB, St-Cyr A, A guide to RBF-generated finite differences for nonlinear transport: shallow water simulations on a sphere, *J. Comput. Phys.* 231 (2012) 4078–4095.
- [11]. Piret C, The orthogonal gradients method: A radial basis functions method for solving partial differential equations on arbitrary surfaces, *J. Comput. Phys.* 231 (20) (2012) 4662–4675.
- [12]. Piret C, Dunn J, Fast rbf ogr for solving pdes on arbitrary surfaces, *AIP Conference Proceedings* 1776 (1).
- [13]. Fuselier EJ, Wright GB, A high-order kernel method for diffusion and reaction-diffusion equations on surfaces, *J. Sci. Comput.* 56 (3) (2013) 535–565.
- [14]. Shankar V, Wright GB, Kirby RM, Fogelson AL, A radial basis function (RBF)-finite difference (FD) method for diffusion and reaction-diffusion equations on surfaces, *J. Sci. Comput.* 63 (3) (2014) 745–768.
- [15]. Lehto E, Shankar V, Wright GB, A radial basis function (RBF) compact finite difference (FD) scheme for reaction-diffusion equations on surfaces, *SIAM J. Sci. Comput.* 39 (2017) A2129–A2151.
- [16]. Fasshauer GE, *Meshfree Approximation Methods with MATLAB*, Interdisciplinary Mathematical Sciences - Vol. 6, World Scientific Publishers, Singapore, 2007.
- [17]. Fornberg B, Wright G, Stable computation of multiquadric interpolants for all values of the shape parameter, *Comput. Math. Appl.* 48 (2004) 853–867.
- [18]. Fornberg B, Piret C, A stable algorithm for flat radial basis functions on a sphere, *SIAM J. Sci. Comput.* 30 (2007) 60–80.
- [19]. Fasshauer GE, McCourt MJ, Stable evaluation of Gaussian radial basis function interpolants, *SIAM J. Sci. Comput.* 34 (2012) A737–A762.
- [20]. Fornberg B, Larsson E, Flyer N, Stable computations with Gaussian radial basis functions, *SIAM J. Sci. Comput.* 33(2) (2011) 869–892.
- [21]. Fornberg B, Lehto E, Powell C, Stable calculation of Gaussian-based RBF-FD stencils, *Comput. Math. Appl.* 65 (2013) 627–637.
- [22]. Wright GB, Fornberg B, Stable computations with flat radial basis functions using vector-valued rational approximations, *J. Comput. Phys.* 331 (2017) 137–156.

- [23]. Bayona V, Flyer N, Fornberg B, Barnett GA, On the role of polynomials in RBF-FD approximations: II. Numerical solution of elliptic PDEs, *J. Comput. Phys.* 332 (2017) 257–273.
- [24]. Shankar V, The overlapped radial basis function-finite difference (RBF-FD) method: A generalization of RBF-FD, *J. Comput. Phys.* 342 (2017) 211–228.
- [25]. Ma H, Chebyshev–legendre super spectral viscosity method for nonlinear conservation laws, *SIAM J. Numer. Anal.* 35 (3) (1998) 893–908.
- [26]. Tadmor E, Convergence of spectral methods for nonlinear conservation laws, *SIAM J. Numer. Anal.* 26 (1) (1989) 30–44.
- [27]. Ma H, Chebyshev–legendre spectral viscosity method for nonlinear conservation laws, *SIAM J. Numer. Anal.* 35 (3) (1998) 869–892.
- [28]. Wendland H, Scattered data approximation, Vol. 17 of Cambridge Monogr. Appl. Comput. Math, Cambridge University Press, Cambridge, 2005.
- [29]. Davydov O, Schaback R, Minimal numerical differentiation formulas, submitted (2017).
- [30]. Behrens J, Iske A, Grid-free adaptive semi-Lagrangian advection using radial basis functions, *Comput. Math. Appl.* 43 (3) (2002) 319–327.
- [31]. Shankar V, Wright GB, Mesh-free semi-lagrangian methods for transport on a sphere using radial basis functions, submitted (2017).
- [32]. Shankar V, Narayan A, Kirby RM, RBF-LOI: Augmenting radial basis functions (RBFs) with least orthogonal interpolation (LOI) for solving PDEs on surfaces, submitted (2017).
- [33]. Shankar V, Kirby RM, Fogelson AL, Robust node generation for meshfree discretizations on irregular domains and surfaces, submitted (2017).
- [34]. Platte RB, Driscoll TA, Eigenvalue Stability of Radial Basis Function Discretizations for Time-dependent Problems, *Comput. Math. Appl.* 51 (8) (2006) 1251–1268.
- [35]. Ascher UM, Ruuth SJ, Wetton BTR, Implicit-Explicit Methods For Time-Dependent PDEs, *SIAM J. Numer. Anal.* 32 (1997) 797–823.
- [36]. Fornberg B, Flyer N, Fast generation of 2-D node distributions for mesh-free PDE discretizations, *Comput. Math. Appl.* 69 (7) (2015) 531–544.
- [37]. Leiderman KM, Fogelson AL, Grow with the flow: a spatial-temporal model of platelet deposition and blood coagulation under flow, *Math. Med. Biol.* 28 (2011) 47–84. [PubMed: 20439306]
- [38]. Leiderman K, Fogelson AL, An Overview of Mathematical Modeling of Thrombus Formation Under Flow, *Thromb. Res.* 133 Suppl (2014) S12–S14. [PubMed: 24759131]
- [39]. Strang G, On the construction and comparison of difference schemes, *SIAM J. Numer. Anal.* 5 (3) (1968) 506–517.
- [40]. Marchuk GI, Some application of splitting-up methods to the solution of mathematical physics problems, *Aplikace matematiky* 13 (2) (1968) 103–132.
- [41]. Saad Y, Schultz MH, Gmres: A generalized minimal residual algorithm for solving nonsymmetric linear systems, *SIAM Journal on Scientific and Statistical Computing* 7 (3) (1986) 856–869.
- [42]. Saad Y, Iterative Methods for Sparse Linear Systems, SIAM, Philadelphia, 2003.

Highlights

- We present a new hyperviscosity-based stabilization technique for RBF-FD discretizations.
- The amount of hyperviscosity is determined quasi-analytically.
- Our new technique is free of tuning parameters.
- We validate this technique to advection-diffusion equations on irregular domains.
- Our method shows high orders of convergence over a range of Peclet numbers.
- We apply this technique to a nonlinearly-coupled transport problem in the spherical shell.

Figure 1:

Spectrum of the discrete Laplacian L on a node set with $N = 658$ nodes for different relationships between m and ℓ . The node set here was a set of Poisson disk samples generated on the disk using the method from [33].

Figure 2:

Spectrum of the discrete Laplacian L on a node set with $N = 919$ nodes for different relationships between m and ℓ . The node set here was a set of Poisson disk samples generated on the disk using the method from [33].

Figure 3:

Local Δ -Lebesgue functions for different scaling laws on a logarithmic scale. The figure contrasts (18) (left), (19) (middle), and (20) (right). Lighter colors indicate larger values, which in turn indicate greater potential for instability.

Figure 4:

Node sets on the disk and in the ball. The above node sets were generated using the algorithms in [33]. Interior nodes are shown as filled circles, and boundary nodes as empty circles.

Figure 5: Relative ℓ_2 error vs $\sqrt{N}$ as a function of stencil size n and polynomial degree ℓ for advection on the unit disk. The figure on the left shows results without filtering, and the figure on the right with filtering. The dashed lines are lines of best fit indicating the slope (and hence convergence rate).

Figure 6:

Relative ℓ_2 error vs $\sqrt[3]{N}$ as a function of stencil size n and polynomial degree ℓ for forced advection-diffusion on the unit disk. The dashed lines are lines of best fit indicating the slope (and hence convergence rate).

Figure 7:

Relative ℓ_2 error vs $\sqrt[3]{N}$ as a function of stencil size n and polynomial degree ℓ . The dashed lines are lines of best fit indicating the slope (and hence convergence rate).

Figure 8.

Relative errors vs node spacing as a function of stencil size n and polynomial degree ℓ for both c (left) and C_B (right).

Table 1

Table of parameters based on desired approximation order ξ , differential operator order θ , node spacing h , velocity $\mathbf{u}$, diffusion coefficient ν , and dimension d .

Parameter	Meaning	Value
ℓ	Polynomial degree for non-hyperviscosity terms	$\xi + \theta - 1$
m	PHS degree	If ℓ is odd, $\ell - 1$ if ℓ is even
M	Number of polynomial terms	$\binom{\ell + d}{d}$
n	Stencil size	$2M + \lfloor \ln(2M) \rfloor$
δ	Overlap parameter	0.7 if $\ell \leq 3$, 0.5 if $4 \leq \ell < 6$, 0.3 if $\ell \geq 6$
k	Order of hyperviscosity	$\lfloor 1.5 \ln n \rfloor$
γ	Magnitude of hyperviscosity	$(-1)^{1-k} 2^{q_1-2k} h^{2k-q_1} \ \mathbf{u}\ _{\infty} \mathbf{e} + (-1)^{1-k} 2^{q_2-2k} h^{2k-q_2} \nu \boldsymbol{\eta}$
$\phi_{\text{hyp}}(r)$	PHS RBF for hyperviscosity	r^{2k+1}
$\ell_{\text{hyp}}(r)$	Polynomial degree for hyperviscosity	k
n_{hyp}	Stencil size for hyperviscosity	$2 \binom{\ell_{\text{hyp}} + d}{d} + 1$