

On the role of polynomials in RBF-FD approximations: II. Numerical solution of elliptic PDEs

Victor Bayona *

Department of Mathematical Sciences
University of Bath
Bath BA2 7AY, UK

Natasha Flyer †

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

Bengt Fornberg ‡ and Gregory A. Barnett §

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

September 1, 2016

Abstract

RBF-generated finite differences (RBF-FD) have in the last decade emerged as a very powerful and flexible numerical approach for solving a wide range of PDEs. We find in the present study that combining polyharmonic splines (PHS) with multivariate polynomials offers an outstanding combination of simplicity, accuracy, and geometric flexibility when solving elliptic equations in irregular (or regular) regions. In particular, the drawbacks on accuracy and stability due to Runge's phenomenon are overcome once the RBF stencils exceed a certain size due to an underlying minimization property. Test problems include the classical 2-D driven cavity, and also a 3-D global electric circuit problem with the earth's irregular topography as its bottom boundary. The results we find are fully consistent with previous results for data interpolation.

* *Email:* victor.bayona.revilla@gmail.com

† *Email:* flyer@ucar.edu

‡ *Email:* fornberg@colorado.edu

§ *Email:* gregory.barnett@colorado.edu

1 Introduction

In this follow-up paper to our previous work in [10], which focused on interpolation, we address here the use of RBF-FD approximations based on polyharmonic splines (PHS) with polynomial augmentation for the numerical solution of elliptic PDEs. This approach benefits from the properties of RBF-FD methods (such as high-order convergence, local node refinement, geometric flexibility, ease of program, etc) while it bypasses complications related to the selection of a shape parameter, ill-conditioning of some RBF collocation matrices and potential convergence failure under refinement (also known as stagnation errors) [9, 10, 11, 14, 15]. We further find that the present approach is able to circumvent the stability issues due to the Runge phenomenon when the stencil size exceeds a certain size for a given fixed polynomial degree.

RBF-FD approximations based on PHS with polynomial augmentation are introduced in Section 2. Section 3 describes some concepts used throughout the paper, such as the ghost node approach for handling boundaries, the algorithm used to discretize the domains and how the linear systems of equations are solved. The main results of this study are contained in Section 4, where we discuss the use of RBF-FD approximations based on PHS with polynomial augmentation for the solution of elliptic PDE problems. In Section 5 the method is applied to solve elliptic PDE problems of increasing complexity. These include the 2-D lid driven cavity and the 3-D Global Electric Circuit. The conclusions are given in Section 6.

2 RBF-FD formulation

RBF-FD methods can be viewed as a natural generalization of FD methods. As in FD, RBF-FD approximates a linear operator L at a location \underline{x}_c as a linear combination of the values of the function at the n closest nodes which form the stencil,

$$Lu|_{\underline{x}_c} \approx \sum_{i=1}^n w_i u_i. \tag{1}$$

In classical FD, the weights w_i are determined by enforcing the approximation (1) to be exact for polynomials of as high degrees as possible. However, this approach is grid based and does not generalize well to scattered nodes in more than 1-D, with the Mairhuber-Curtis theorem being just one reason. Instead of only relying on multivariate polynomials, we include also RBFs that are centered at the nodes of the stencil of size n . This approach for finding weights is guaranteed to be non-singular no matter how the n nodes are scattered in any number of dimensions (assuming the nodes are distinct and a unisolvency condition is met).

2.1 RBF-FD weights

When not augmenting the RBFs with polynomials, the RBF-FD weights w_i in (1) based on stencil nodes $\underline{x}_i, i = 1, 2, \dots, n$ are obtained by solving

$$\begin{bmatrix} \phi(\|\underline{x}_1 - \underline{x}_1\|) & \dots & \phi(\|\underline{x}_1 - \underline{x}_n\|) \\ \vdots & \ddots & \vdots \\ \phi(\|\underline{x}_n - \underline{x}_1\|) & \dots & \phi(\|\underline{x}_n - \underline{x}_n\|) \end{bmatrix} \begin{bmatrix} w_1 \\ \vdots \\ w_n \end{bmatrix} = \begin{bmatrix} L\phi(\|\underline{x} - \underline{x}_1\|)|_{\underline{x}_c} \\ \vdots \\ L\phi(\|\underline{x} - \underline{x}_n\|)|_{\underline{x}_c} \end{bmatrix}. \tag{2}$$

There are many choices of RBFs $\phi(r)$ that can be used [7, 14]. In this work, we limit our attention to the use of polyharmonic splines (PHS, here $\phi(r) = r^{2m-1}, m \in \mathbb{N}$) with polynomial augmentation

[9, 10, 17]. For example, using RBF-FD + linear terms, the RBF-FD weights w_i in 2-D are obtained by solving

$$\begin{bmatrix} & & & 1 & x_1 & y_1 \\ & & & \vdots & \vdots & \vdots \\ & & & 1 & x_n & y_n \\ \hline & A & & & & \\ 1 & \dots & 1 & & & \\ x_1 & \dots & x_n & & 0 & \\ y_1 & \dots & y_n & & & \end{bmatrix} \begin{bmatrix} w_1 \\ \vdots \\ w_n \\ \hline \gamma_1 \\ \gamma_2 \\ \gamma_3 \end{bmatrix} = \begin{bmatrix} L\phi(\|\underline{x} - \underline{x}_1\|)|_{\underline{x}_c} \\ \vdots \\ L\phi(\|\underline{x} - \underline{x}_n\|)|_{\underline{x}_c} \\ \hline L1|_{\underline{x}_c} \\ Lx|_{\underline{x}_c} \\ Ly|_{\underline{x}_c} \end{bmatrix}, \quad (3)$$

where A is the same matrix as in (2). More generally, the linear system of equations takes the form

$$\begin{bmatrix} A & P \\ P^T & 0 \end{bmatrix} \begin{bmatrix} \mathbf{w} \\ \boldsymbol{\gamma} \end{bmatrix} = \begin{bmatrix} L\boldsymbol{\phi} \\ L\mathbf{p} \end{bmatrix}, \quad (4)$$

where \mathbf{w} and $L\boldsymbol{\phi}$ are $n \times 1$ column vectors formed by w_i and $L\phi(\|\underline{x} - \underline{x}_i\|)|_{\underline{x}=\underline{x}_c}$, P^T is a $m \times n$ matrix formed by all the $m = \binom{p+d}{p}$ multivariate polynomial terms of total degree less than or equal to p in d dimensions (c.f. Table 1), and $\boldsymbol{\gamma}$ and $L\mathbf{p}$ are $m \times 1$ vectors formed by γ_i and $Lp_i(\underline{x})|_{\underline{x}=\underline{x}_c}$. The scalar p denotes the highest degree of the polynomials $p_i(\underline{x})$, $i = 1, 2, \dots, m$. Notice that a necessary but not sufficient condition for unisolvency is that $n \leq m$ (for a rigorous discussion on unisolvency conditions see [7]).

dimension	Polynomial degree p								
d	0	1	2	3	4	5	6	7	...
1	1	2	3	4	5	6	7	8	...
2	1	3	6	10	15	21	28	36	...
3	1	4	10	20	35	56	84	120	...

Table 1: The number of polynomial terms m up through degree p , when in d dimensions; $m = \binom{p+d}{p}$.

2.2 Role of polynomials

As previously shown in the first part of our work [10], the solution of the linear system (4) is the solution of the equality-constrained quadratic programming problem

$$\min_{\mathbf{w}} J(\mathbf{w}) = \frac{1}{2} \mathbf{w}^T A \mathbf{w} - \mathbf{w}^T L\boldsymbol{\phi} \quad \text{subject to} \quad P^T \mathbf{w} = L\mathbf{p}, \quad (5)$$

This result indicates that the RBF-FD weights \mathbf{w} are the solution of the under-determined linear system $P^T \mathbf{w} = L\mathbf{p}$ when minimizing the functional $J(\mathbf{w})$. In Sections 4, we will see that allowing this system to become more undetermined by increasing the stencil size such that $n \gtrsim 2m$ leads to optimal RBF-FD weights in the context of solving elliptic PDEs.

3 Some preliminaries

3.1 Ghost nodes

The ghost (or fictitious) node approach has been used frequently in the past to implement boundary conditions both for finite differences [22] and pseudospectral methods [5, 12]. The underlying idea is to avoid the use of fully one-sided stencils in order to improve accuracy and stability. This approach was described in [8, 18] for solving elliptic PDEs with global RBFs. In the present work, the ghost node approach is used in the context of RBF-FD approximations.

Consider an elliptic PDE problem defined in a domain Ω . To avoid one-sided stencils near the boundary, a layer of ghost nodes with as many nodes as boundary nodes is introduced outside the domain all around the boundary. Using these nodes to discretize the PDE avoids fully one-sided stencils. However, the resulting $N \times (N + N_g)$ differentiation matrix is not longer square, where N is the number of interior and boundary nodes and N_g is the number of ghost nodes. In this case, the differentiation matrix can be written as

$$\begin{bmatrix} D & D_g \end{bmatrix} \begin{bmatrix} \underline{u} \\ \underline{u}_g \end{bmatrix} = \underline{f}, \quad (6)$$

where the block matrices are $(D)_{N \times N}$ and $(D_g)_{N \times N_g}$, \underline{u} is the value of the unknowns on the interior and boundary nodes, \underline{u}_g is the value of the unknowns on the ghost nodes and \underline{f} is the right-hand side evaluated on the interior nodes.

In order to have a linear system of equations with as many unknowns as conditions, the PDE can also be enforced on the boundary, which leads to the square differentiation matrix,

$$\begin{bmatrix} D & D_g \\ B & B_g \end{bmatrix} \begin{bmatrix} \underline{u} \\ \underline{u}_g \end{bmatrix} = \begin{bmatrix} \underline{f} \\ \underline{f}_g \end{bmatrix}, \quad (7)$$

where $(B)_{N_g \times N}$ and $(B_g)_{N_g \times N_g}$ are the matrices that contain the weights discretizing the PDE on the boundary and \underline{f}_g is the right-hand side evaluated on the boundary nodes. By simple elimination, the differentiation matrix can be written as a function of the interior and boundary values as

$$\underbrace{[D - D_g B_g^{-1} B]}_{\hat{D}} \underline{u} = \underbrace{[\underline{f} - D_g B_g^{-1} \underline{f}_g]}_{\hat{\underline{f}}}. \quad (8)$$

For large stencil sizes, when a single layer of ghost nodes might not be enough to relieve the boundary effects, it is possible to increase the number of layers. For each layer, a new condition must be enforced on the boundary. A simple way to obtain these extra conditions for linear PDEs is by applying powers of the PDE. For example, $Lu = f$ for the first layer, $L^2u = Lf$ for the second layer, $L^3u = L^2f$ for the third layer, and so on (assuming that f is of a form such that we can evaluate Lf , L^2f , etc)

3.2 Domain discretization

The domains Ω considered in this work are discretized using scattered nodes based on the following repel algorithm. Other options for creating 2-D node distribution are described in [13, 20].

1. Discretize the boundary $\partial\Omega$ with equispaced nodes. Root-finding algorithms might be useful for this purpose.
2. Embed the discretized domain Ω into a regular domain (e.g. a rectangle in 2-D) and discretize it using structured nodes. The size of this regular domain must be at least 4 layers of nodes larger all around the embedded domain. The 2-3 most exterior layers of nodes are called frame nodes.
3. Discard the nodes lying within a certain distance to the boundary nodes.
4. Boundary and frame nodes are fixed to their positions. A slightly random displacement is applied to the rest of the nodes, followed by calculating the repulsion force

$$\vec{F} = \sum_{i=1}^n \frac{\vec{r}_i}{\|\vec{r}_i\|_2^3},$$

where $\vec{r}_i = (x - x_i, y - y_i)$ and n is the n -th closest node to the i -th node. Typically only the direction of \vec{F} is utilized, with actual node movement distances decreased in a systematic manner during a small number of iterations on the order of 20 or so.

5. The domain Ω is finally discretized using the interior and boundary nodes. All nodes outside Ω are discarded. In the case that ghost nodes are used, they are separated placed and simply frame the boundary as illustrated by the crosses in Figure 1.

3.3 Elliptic solver and preconditioner

Unless specified otherwise in the text, the linear system of equations is solved using the BiCGSTAB iterative method with incomplete LU factorization as preconditioner (with 0 level of fill in) and reverse Cuthill-McKee ordering. The residual tolerance is set to 10^{-11} and the maximum number of iterations is equal to 100.

4 PHS RBF-FD approximations with polynomial augmentation for the numerical solution of elliptic PDEs

4.1 A test problem to explore accuracy

In the following test problem, we explore the performance of PHS RBF-FD with polynomial augmentation when solving Poisson's equation

$$\begin{cases} \Delta u = f(x, y), & \text{in } \Omega \\ u = g(x, y), & \text{on } \partial\Omega \end{cases} \quad (9)$$

on the unit disc $\Omega = \{(x, y), x^2 + y^2 \leq 1\}$, where $f(x, y)$ and $g(x, y)$ are computed from the test function

$$u(x, y) = \sin[10(x + y)]. \quad (10)$$

Scattered nodes based on the repel algorithm from Section 3.2 are used to discretize the domain. Figure 1 shows the node distribution for $h \approx 0.05$, where the dots, squares and crosses represent interior, boundary and exterior ghost nodes, respectively. The node distributions for $h = 0.025$ and $h = 0.01$ are structurally similar, although denser.




Figure 1: Unit disc discretized using $h \approx 0.05$: interior nodes (dot), boundary nodes (square) and exterior nodes (cross). Exterior nodes are treated differently depending on the approach used to handle the boundary.

As it is well-known, one-sided stencils near boundaries can affect the accuracy and stability of elliptic solvers. To explore the effect of boundaries on PHS RBF-FD approximations with polynomials, we solve (9) using three different approaches to handle the boundary. Figure 2 shows the relative ℓ_2 norm error as a function of the stencil size and polynomial degree for different resolutions (row of subplots) and methods for handling the boundary (column of subplots). The dashed curves $n = \binom{p+2}{p}$ plotted in each figure represent the case where the stencil size n is equal to the number polynomial terms m (and thus $P^T \mathbf{w} = L\mathbf{p}$ is square and no minimization of $J(\mathbf{w})$ can occur). For higher values of n (above the dashed curves), the collocation matrix (4) is generally non-singular; below it, the matrix will be singular since the columns of P will be linearly dependent. The dashed-dot curves, representing $n = 2 \cdot \binom{p+2}{p}$ in the right column of subplots, will be explained below.

In the first column of subplots, we enforce the analytical solution (10) on both boundary and the three layers of exterior nodes. This is an ideal case, leading to an optimal reference solution. Observe that the error decreases as the node distribution is refined and also as the polynomial order increases. On the contrary, increasing the stencil size does not significantly improve the accuracy, especially for fine node distributions. Note also that only for $h = 0.01$ does the accuracy deteriorate when the stencil size n is very close to the smallest stencil size that can be used for a given polynomial degree.

The second column of subplots shows the result of using the ghost node approach explained in Section 3.1 to handle the boundary. In this case, the three layers of exterior nodes are used to discretize the equation and, in order to have a linear system of equations with as many unknowns as equations, the PDE $\Delta u = f$, together with $\Delta^2 u = \Delta f$ and $\Delta^3 u = \Delta^2 f$, are enforced on the boundary. As a result, fully one-sided stencils are avoided and the accuracy and stability of the solver are not compromised. The results here are very similar to the ideal ones from the first column, which makes this approach an effective method to avoid boundary effects in PDE problems.

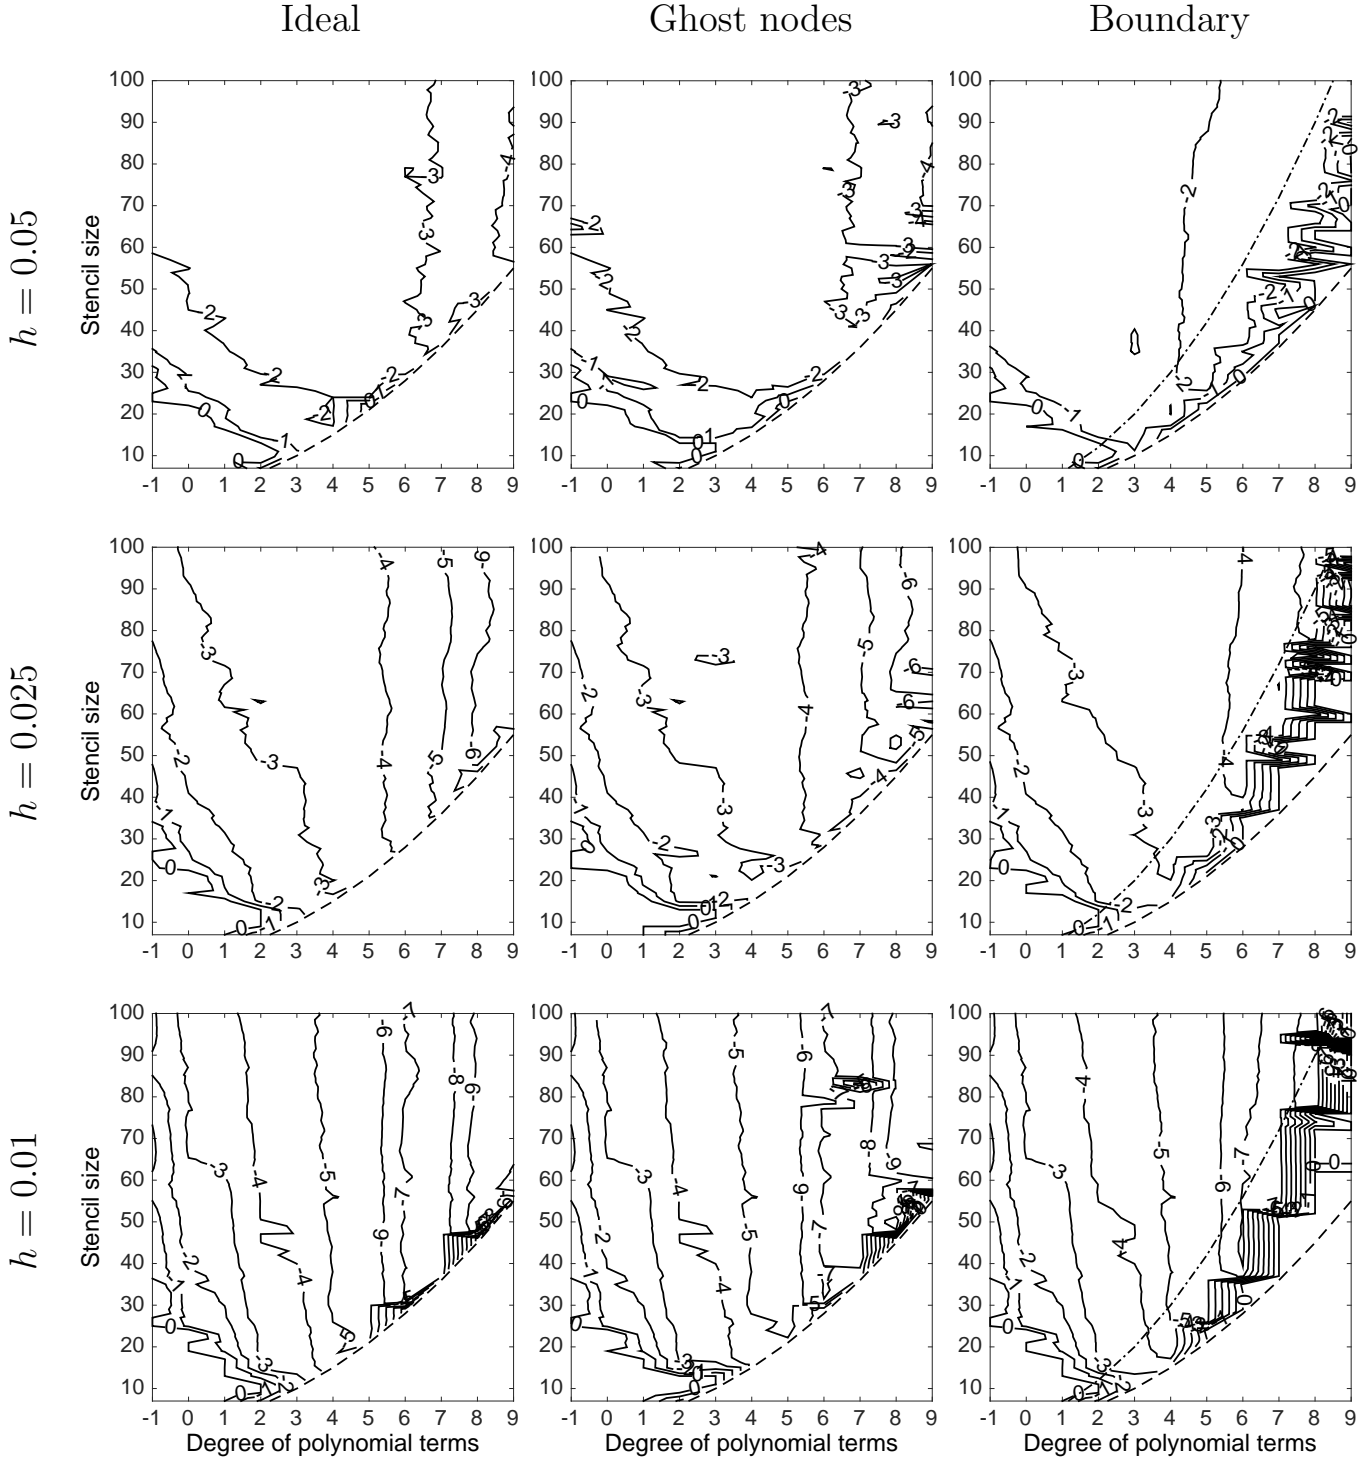


Figure 2: Accuracy (\log_{10} relative ℓ_2 norm error) when using PHS r^7 with polynomial augmentation to solve (9) as a function of the stencil size and degree of included polynomial terms. The rows show different resolutions and the columns different approaches to handle the boundaries. The dashed line shows the smallest stencil size $n = \binom{p+2}{p}$ that can be used for each degree of polynomial term. The dashed-dot line on the right column of subplots shows twice the number of polynomial terms $n = 2 \cdot \binom{p+2}{p}$.

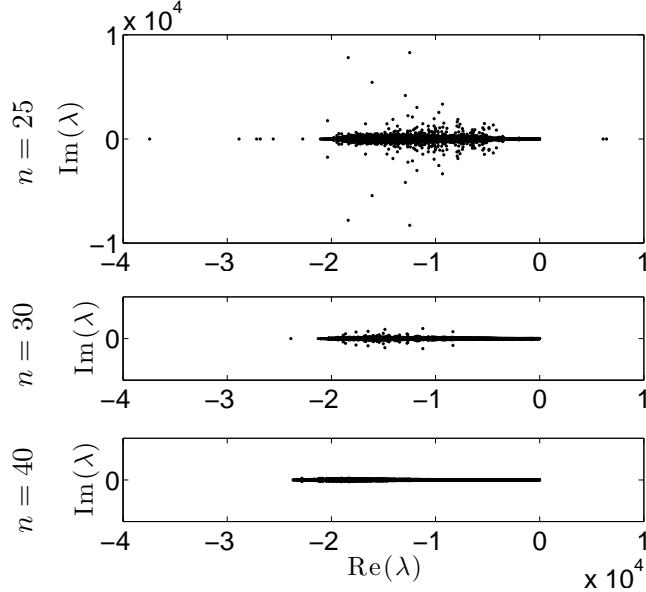


Figure 3: Eigenvalue distribution of the differentiation matrix that approximates (9) for the boundary case with $h \approx 0.025$, 5th degree polynomial terms and $n = 25, 30$ and 40 . The scales on the axis are the same in the three subplots.

The third column shows the result of solving the problem directly without any special treatment of the boundary. In this case, exterior nodes are never introduced and only interior and boundary nodes are considered to discretize the PDE and the boundary conditions. As a result, it leads to one-sided stencils near the boundary. Under refinement, the boundary effects are seen to be significant only for stencil sizes smaller than twice the number of polynomial terms included, that is below the dashed-dot curves in the subplots. For larger stencil sizes, it remarkably reproduces the same results as the ideal case. This is consistent with test results in [18] (using smooth global RBFs), which showed only insignificant differences between using ghost nodes and not doing so, with no special treatment of the boundaries.

In this case, the eigenvalue distribution of the differentiation matrix (i.e. the Laplacian) also reflects this behaviour as a function of the stencil size. Figure 3 shows the eigenvalue distribution for $h \approx 0.025$, with polynomial terms up to 5th degree and $n = 25, 30$ and 40 . Notice that for $n = 25$, the eigenvalues are widely scattered over the complex plane, some of them with positive real part. As the stencil size increases, the eigenvalue distribution improves, becoming closer to the negative real axis. In the particular case $n = 40$, which corresponds to almost twice as many nodes as polynomials, the eigenvalues lie very close to the negative real axis, as expected for a differentiation matrix approximating the Laplacian.

4.2 An observation about PHS RBF-FD accuracy near boundaries

When not using ghost nodes outside the boundary, accuracy and stability turn out to become excellent once the RBF stencils exceed a certain size. This observation is potentially very useful as it may make it unnecessary to have to deal with ghost points, but counter intuitive, since larger stencils might be expected to suffer more from near-one-sidedness than smaller ones. The purpose of this 1-D test case is to illustrate what seems to be the mechanism behind this advantageous

1.0000	-2.0000	1.0000									
1.0000	-2.0000	1.0000	0								
0.9167	-1.6667	0.5000	0.3333	-0.0833							
0.8333	-1.2500	-0.3333	1.1667	-0.5000	0.0833						
0.7611	-0.8167	-1.4167	2.6111	-1.5833	0.5167	-0.0722					
0.7000	-0.3889	-2.7000	4.7500	-3.7222	1.8000	-0.5000	0.0611				
0.6482	0.0254	-4.1500	7.6500	-7.3472	4.7000	-1.9500	0.4754	-0.0518			
0.6040	0.4236	-5.7429	11.3667	-12.9222	10.2750	-5.6667	2.0683	-0.4500	0.0442		

Table 2: Standard FD weights (rows) approximating $u''(x)$ at $x = 0$ over the points $x = -1, 0, 1, 2, \dots, n-2$ (columns) for increasing degrees $p = 2$ (first row), $p = 3$ (second row), $\dots, p = 9$ (last row) of polynomials and stencil sizes $n = p + 1$.

0.7000	-0.3889	-2.7000	4.7500	-3.7222	1.8000	-0.5000	0.0611										
0.7591	-0.8617	-1.0450	1.4400	0.4152	-1.5100	1.1550	-0.4117	0.0591									
0.8081	-1.1947	-0.1461	0.3524	0.5318	-0.1117	-0.7871	0.8534	-0.3654	0.0592								
0.8513	-1.4552	0.4419	-0.1590	0.4558	0.0910	-0.1708	-0.4492	0.6677	-0.3332	0.0597							
0.8901	-1.6661	0.8493	-0.4110	0.3399	0.1458	-0.0158	-0.1595	-0.2643	0.5394	-0.3079	0.0602						
0.9252	-1.8416	1.1460	-0.5448	0.2554	0.1197	0.0713	-0.0712	-0.1199	-0.1605	0.4468	-0.2870	0.0606					
0.9572	-1.9901	1.3696	-0.6178	0.1977	0.0817	0.0913	0.0023	-0.0761	-0.0866	-0.0976	0.3766	-0.2688	0.0607				
0.9866	-2.1177	1.5426	-0.6582	0.1606	0.0453	0.0889	0.0318	-0.0211	-0.0708	-0.0591	-0.0589	0.3220	-0.2528	0.0606			
1.0137	-2.2283	1.6794	-0.6803	0.1376	0.0154	0.0772	0.0427	0.0073	-0.0323	-0.0599	-0.0386	-0.0343	0.2785	-0.2383	0.0603		

Table 3: Cubic PHS RBF-FD weights (rows) augmented with polynomials up to 7th degree, approximating $u''(x)$ at $x = 0$ over the points $x = -1, 0, 1, 2, \dots, n-2$ (columns) for increasing stencil sizes $n = 8, 9, \dots, 16$ (rows).

feature.

When solving $u''(x) = f(x)$ in 1-D without any ghost nodes, the first location at which $u''(x)$ needs to be approximated is one step in from the boundary. For simplicity, we let the grid spacing be $h = 1$, and consider approximations to $u''(x)$ that are accurate at $x = 0$ for a stencil extending over the points $x = -1, 0, 1, 2, 3, \dots, n-2$. The standard FD weights are readily calculated for increasing values of n (i.e. for increasing degrees $p = n - 1$ of the polynomials used in deriving them) as listed in Table 2.

As the order is increased, the weights near the stencil center grows exponentially fast in magnitude, while the ones in the region of interest (i.e. mainly the leftmost three columns) become less significant. For instance, focusing on the $p = 7$ case (third to last row in Table 2), we next keep this p -value fixed but increase the stencil size from $n = 8$ to $n = 16$. The RBFs will then come to play an increasing role in the process of determining the stencil weights while the formal accuracy order is maintained, as determined by the degree $p = 7$ polynomials. Using for example cubic RBFs, we then obtain the weight sequence in Table 3 for approximating $u''(x)$ at $x = 0$.

We recognize the top row here as row 6 in Table 2 - corresponding to using polynomials only. The following rows show that, when including increasingly many RBFs (while maintaining the same p value), the pattern of weights quickly reverts itself to the ideal case of the weights at the leftmost nodes dominating over all others. By somewhat resembling the pattern $[-1, 2, -1]$ at the three leftmost nodes, $x = -1, 0, 1$, and being smaller everywhere else, they offer similar stability features as the second-order $n = 3$ classical finite difference approximation, yet at a much higher accuracy.

A natural way to test the accuracy that the stencils in Table 3 provide for $u''(x)$ is to apply each of the stencils to the test functions $u_\omega(x) = e^{i\omega x}$, and plot the error (in magnitude) compared to the analytical result $u''_\omega(x) = -\omega^2 e^{i\omega x}$ at the node location of interest (second from the left), for different ω . This gives the results shown in Figure 4. All the curves show the same slope, reflecting that they all are exact for polynomials up through degree $p = 7$. As n is increased, the near-straight-line curves move up slightly.




Figure 4: Absolute error approximating $u''(x)$, compared to the analytical result $u''_{\omega}(x) = -\omega^2 e^{i\omega x}$ at the node location $x = 0$, second from the left, for different ω .

The main conclusion becomes that, purely from the perspective of derivative accuracy (with given data), the approximation using only polynomials and no RBFs is the most accurate option. However, this version is unacceptable in the context of solving an elliptic equation due to stability issues. As we include more nodes (increase n), the errors are seen to increase only slowly, making the somewhat larger n -cases good compromises between accuracy and providing near-diagonally dominant stencils. In the larger context of solving elliptic systems, this small loss in accuracy is insignificant compared to the gain in numerical stability.

Alternatively, this feature can be approached through cardinal functions. These take the value one at one node point and zero at the remaining ones. In Figure 5a (upper part), we see the $n = 8$ cardinal functions for nodes located at $x = -1, 0, 1, 2, \dots, 6$, when interpolating with pure polynomials of degree $p = 7$. The corresponding weights for approximating $u''(x)$ at $x = 0$ can be read off as the second derivative at $x = 0$ of these successive cardinal functions - displayed in the lower part of Figure 5a. These match the numbers in the top row of Table 3.

The upper parts of Figures 5b and 5c show how the individual cardinal functions become less prone to spurious end oscillations when increasingly many nodes are used, i.e. when the relation $P^T \mathbf{w} = L\mathbf{p}$ in (5) becomes underdetermined. The freedoms this creates in the weights vector \mathbf{w} will then be settled by the minimization part in (5). The larger n is, the more extensive will this minimization become. As a result, the stencil weights decrease in magnitude wherever this is possible, given that the order of accuracy needs to be preserved.


5 Numerical solution of elliptic PDE problems: Additional test cases

5.1 Elliptic PDE with variable coefficients


In this section, we explore further the convergence of PHS with polynomial augmentation as a function of the degree of the polynomial terms. For this purpose, we consider the general elliptic PDE with variable coefficients,

$$\Delta u + \eta(x, y) \frac{\partial^2 u}{\partial x \partial y} + \alpha(x, y) \frac{\partial u}{\partial x} + \beta(x, y) \frac{\partial u}{\partial y} + \gamma(x, y) u = f(x, y), \quad \text{in } \Omega \quad (11)$$

(a) Cardinal functions; $n = 8, p = 7$.



(b) Cardinal functions; $n = 10, p = 7$.



(c) Cardinal functions; $n = 16, p = 7$.




Figure 5: Cardinal functions and associated weights \mathbf{w} when using polynomials of degree $p = 7$ together with cubic PHS and increasing number of stencil nodes n . Since each successive RBF-FD weight is equal to the curvature (i.e. second derivative) of the corresponding cardinal functions near $x = 0$, these segments have been highlighted by a bold line style.




Figure 6: Discretized domain Ω for $h \approx 0.025$ using node repel algorithm from Section 3.

and Dirichlet boundary conditions

$$u = g(x, y), \quad \text{on } \partial\Omega \quad (12)$$

on the domain Ω plotted in Figure 6, where the inner and outer boundaries in polar coordinates are

$$r(\theta) = \frac{3}{10} + \frac{1}{10} \sin(\theta) + \frac{3}{20} \sin(5\theta)$$

and

$$r(\theta) = 1 + \frac{1}{5} \cos(\theta) + \frac{3}{20} \sin(4\theta),$$

respectively. The variable coefficients are given by

$$\begin{aligned} \eta(x, y) &= \cos(10\pi y), \\ \alpha(x, y) &= \exp[-y^2 + \cos(4\pi x) \sin(3\pi y)], \\ \beta(x, y) &= -y \sin(4\pi x), \\ \gamma(x, y) &= x^2 y. \end{aligned}$$

The functions $f(x, y)$ and $g(x, y)$ are obtained assuming the analytical solution

$$u(x, y) = \sin(2\pi y^2 + 3\pi x) - \cos(\pi y - 2\pi x^2). \quad (13)$$

In this test case, the variable coefficients are highly oscillatory and a mixed derivative term $\frac{\partial^2 u}{\partial x \partial y}$, which adversely affects diagonal dominance in FD contexts, is also included. The domain is discretized based on the repel algorithm from Section 3.2. Figure 6 shows the discretized domain for $h \approx 0.025$.

In view of the advantageous feature of PHS RBF-FD near boundaries found in Section 4 (where it was observed that the adverse effects of boundaries on accuracy and stability can be avoided simply by increasing the stencil size), we solve the PDE problem directly without any special treatment of the boundary. This is especially attractive for domains with irregular boundaries, where the ghost nodes approach is not straightforward to apply.

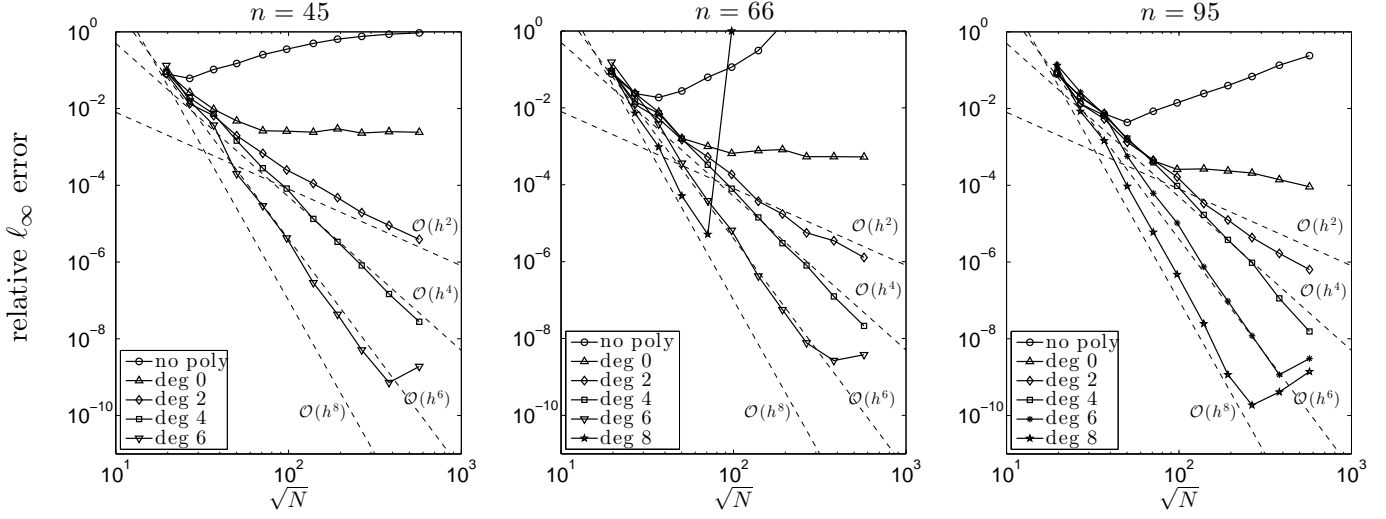


Figure 7: Relative ℓ_∞ error as a function of the number of nodes \sqrt{N} when solving (11) using the stencil sizes $n = 45, 66$ and 95 .

Figure 7 shows the relative ℓ_∞ error as a function of \sqrt{N} for different polynomial degrees and stencil sizes $n = 45, 66$, and 95 . Observe that the convergence order coincides with the degree of polynomials included and is independent of the stencil size n . The stencil size only shifts the curves vertically. Figure 7 also illustrates the effect of increasing the stencil size when boundaries are present. For example, the case $n = 45$ and 8th degree polynomial cannot be computed since the stencil size is equal to the corresponding number of polynomial terms and the unisolvency condition is not met. As the stencil size increases, the constraints imposed on the polynomial basis can be satisfied, leading to the case $n = 95$ (which is about twice the number of polynomial terms) where adverse boundary effects disappear.

Although not explored here, it may be advantageous to use larger stencils only at boundaries (for stability) and smaller inside the domain (for lower cost).

5.2 Non-linear elliptic PDE: Lid driven cavity flow

To test the performance of PHS RBF-FD with polynomial augmentation, we consider in this section the lid-driven cavity flow [1, 3, 16, 19], which is a well-known benchmark problem for both, method assessment and code validation. In the context of RBF-FD methods, it was first considered in [4, 21]. The governing steady-state incompressible Navier-Stokes equations, expressed in the stream function formulation, take the form of the non-linear biharmonic equation

$$\Delta^2 \psi + \text{Re} \left[\frac{\partial \psi}{\partial x} \frac{\partial \Delta \psi}{\partial y} - \frac{\partial \psi}{\partial y} \frac{\partial \Delta \psi}{\partial x} \right] = 0, \quad (14)$$

defined in the unit square $\Omega = [0, 1] \times [0, 1]$. The boundary conditions are

$$\psi = 0, \quad \vec{x} \text{ on } \partial\Omega, \quad (15)$$

and

$$\frac{\partial \psi}{\partial x} \Big|_{(0,y)} = 0, \quad \frac{\partial \psi}{\partial x} \Big|_{(1,y)} = 0, \quad \frac{\partial \psi}{\partial y} \Big|_{(x,0)} = 0, \quad \frac{\partial \psi}{\partial y} \Big|_{(x,1)} = 1. \quad (16)$$

We have used Newton's method to solve this non-linear biharmonic equation. Starting with an initial solution $\psi^{(0)}$, a sequence of approximate solutions $\psi^{(i)}$ are computed as $\psi^{(i)} = \psi^{(i-1)} + \xi$, where the correction ξ is obtained at each iteration by solving the linearized PDE

$$\Delta^2 \xi + \text{Re} \left[\frac{\partial \psi^{(i-1)}}{\partial x} \frac{\partial \Delta \xi}{\partial y} - \frac{\partial \psi^{(i-1)}}{\partial y} \frac{\partial \Delta \xi}{\partial x} + \frac{\partial \xi}{\partial x} \frac{\partial \Delta \psi^{(i-1)}}{\partial y} - \frac{\partial \xi}{\partial y} \frac{\partial \Delta \psi^{(i-1)}}{\partial x} \right] = R(\psi^{(i-1)}), \quad (17)$$

with boundary conditions

$$B\xi = g(x, y) - B\psi^{(i-1)}. \quad (18)$$

In this equation, $R(\psi^{(i-1)})$ is the residual at iteration i , given by

$$R(\psi^{(i-1)}) = \Delta^2 \psi^{(i-1)} + \text{Re} \left[\frac{\partial \psi^{(i-1)}}{\partial x} \frac{\partial \Delta \psi^{(i-1)}}{\partial y} - \frac{\partial \psi^{(i-1)}}{\partial y} \frac{\partial \Delta \psi^{(i-1)}}{\partial x} \right], \quad (19)$$

and B is the operator obtained after writing the boundary conditions (15) and (16) as $B\psi = g(x, y)$. Newton iterations continue until the tolerance $\|R(\psi^{(i-1)})\| \leq \epsilon$ is reached. The algorithm starts with $\text{Re} = 0$ and increases at steps of 100. The procedure above is repeated, starting with the solution from the previous Reynolds number, until the selected Reynolds number has been reached.

To better resolve boundary layers, it is beneficial to use a refined node distribution near the boundary [1, 3]. In this test problem, we have used the node distribution illustrated and described in Figure 8 (with an enlargement of the right bottom corner in the right subplot). The derivatives in (17) are approximated using PHS r^7 augmented with all $\binom{8+2}{8} = 45$ polynomial terms up to 8th degree. The selected stencil size is $n = 90$, which is twice as many nodes as polynomial terms. Since the boundary is of a very simple shape, we choose to use one layer of ghost nodes around the domain to enforce both boundary conditions (15) and (16). The resulting number of unknowns on the boundary are $2N_b + 8$, where N_b is the number of boundary nodes. Enforcing (15) and (16) on the boundary nodes yields to $2N_b + 4$ conditions that, together with the 4 extra conditions obtained by enforcing the PDE (17) also on the corner boundary nodes, form a consistent system of equations.

Figure 9 shows the solution of (14) for $\text{Re} = 0, 1000, 2500, 5000, 7500$ and 10000 . In all the cases, there are three areas where vortices form: at the center of the cavity (Primary vortex), at the bottom-right corner (BR1) and at the bottom-left corner (BL1). The stream function and vorticity values, as well as locations (x, y) of the main vortex at these areas, are computed. These all are listed in Table 4.

The present node distribution, formed by $N = 39,389$ nodes, obtains very competitive results with respect to those found in the literature for $257 \times 257 = 66,049$ nodes [1, 6]. The locations and strengths of the vortices for $\text{Re} = 0$ and 1000 are in very good agreement with those found in [1, 3], while the accuracy seems to deteriorate as the Reynolds number increases. This appears to be a common issue reported frequently in the literature. For example, in [1] accurate results were computed up to $\text{Re} = 12,500$. For $\text{Re} = 16,000$ only qualitative results were shown, concluding that their solution at this Reynolds number was under-resolved and a finer node distribution was needed. In [6] they used a $601 \times 601 = 361,201$ node grid for $\text{Re} = 21,000$.

For qualitative purposes, we have also computed the solutions for $\text{Re} = 15,000$ and $20,000$, but without increasing our previous resolution. The results are plotted in Figure 10 for the same contour values as in Figure 9. Observe that we are also able to resolve the small vortices located




Figure 8: Discretized domain based on a node repel algorithm using $N = 39,389$ nodes: 36,181 interior nodes, 1,600 boundary nodes and 1,608 ghost nodes. The internodal distance increases from $h \approx 0.0025$ near boundaries to $h \approx 0.01$ in the center of the cavity.

at the top-left corner and bottom corners shown in [6] for the case $\text{Re} = 20,000$, but using 10 times less nodes. In order to improve the accuracy of the vortex strengths, an increase of the resolution seems to be also necessary in our case.

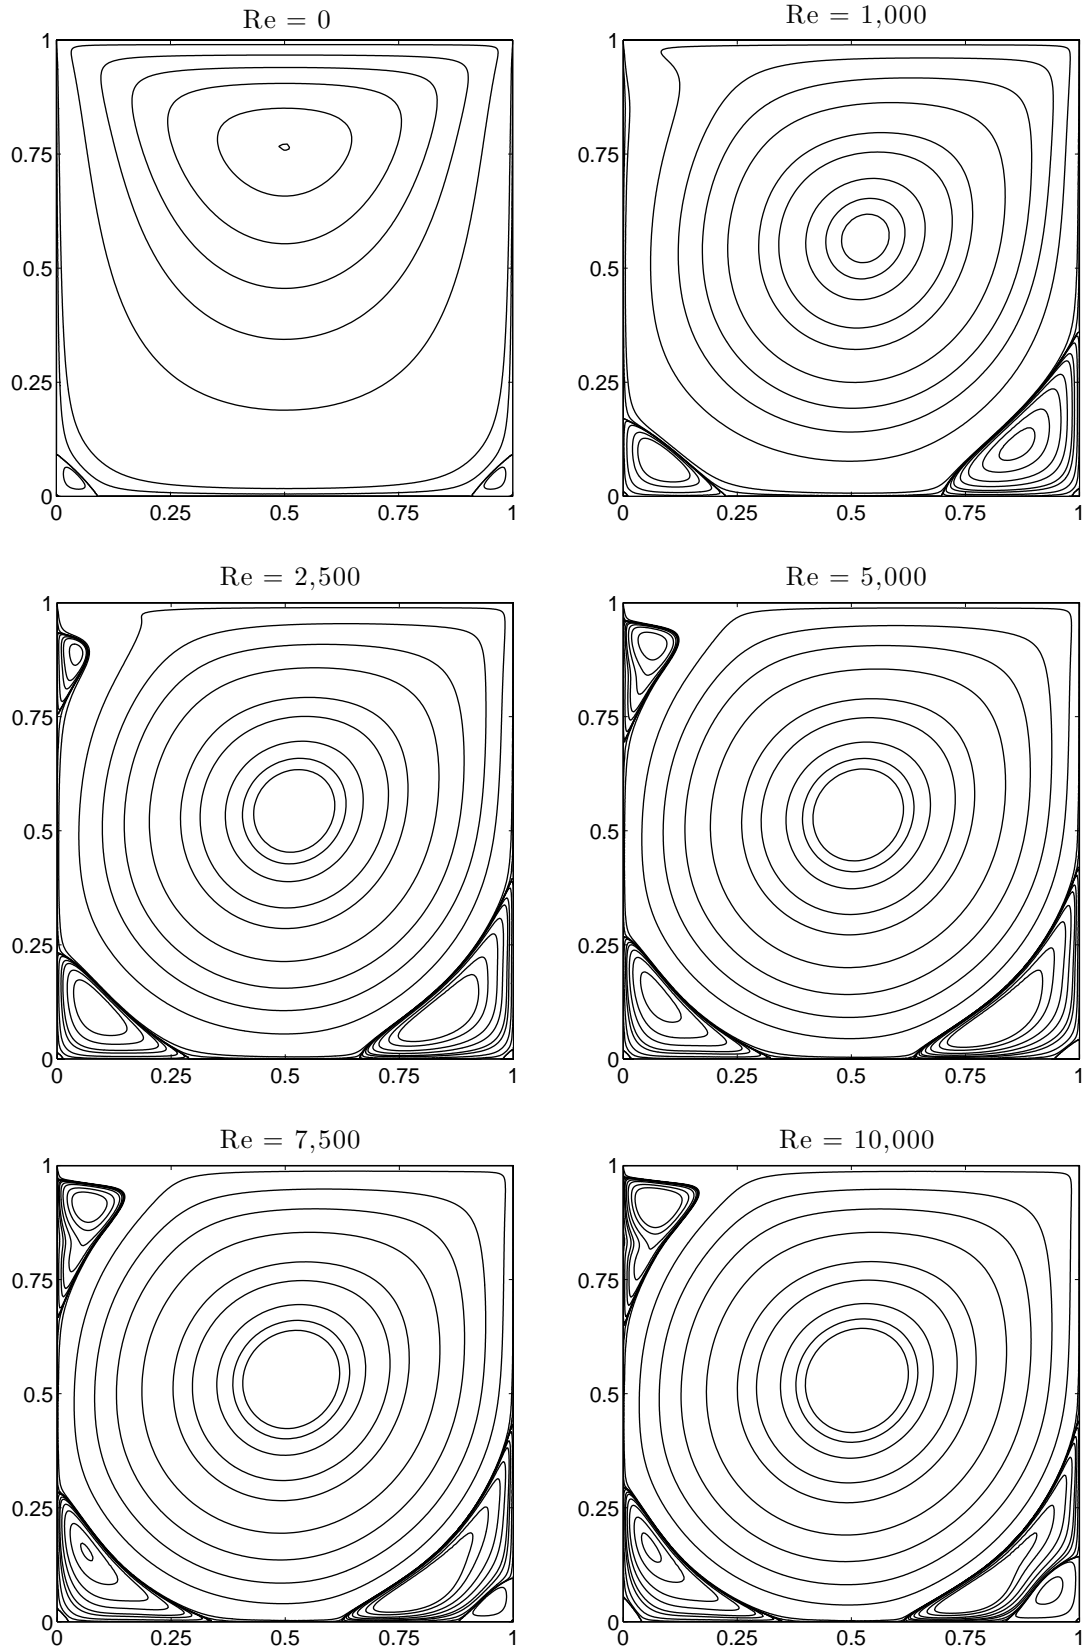


Figure 9: Stream function contours for $Re = 0, 1,000, 2,500, 5,000, 7,500, 10,000$. The contour values are set to $-0.1175, -0.115, -0.11, -0.1, -0.09, -0.07, -0.05, -0.03, -0.01, -1e-4, -1e-5, -1e-10, 0, 1e-6, 1e-5, 5e-5, 1e-4, 2.5e^{-16}, 5e-4, 1e-3$ and $1.5e-3$.




Figure 10: Stream function contours for $Re = 15,000$ (left) and $Re = 20,000$ (right).

Table 4: Stream function and vorticity values at the vortex centers; together with their (x, y) locations.

	Re	0	1,000	2,500	5,000	7,500	10,000
Primary vortex	ψ	-0.1000756	-0.1189307	-0.1214544	-0.1223237	-0.1230018	-0.1237093
	ω	-3.2123303	-2.0676832	-1.9758657	-1.9419612	-1.9360192	-1.9385587
	(x, y)	(0.5000,0.7650)	(0.5308,0.5652)	(0.5198, 0.5439)	(0.5151,0.5352)	(0.5132, 0.5317)	(0.5121,0.5296)
BR1	ψ	0.2227561E-05	0.1729705E-02	0.2661659E-02	0.3077562E-02	0.3240731E-02	0.3216516E-02
	ω	0.0152408	1.1089003	1.9387076	2.7574341	3.2653721	3.8385719
	(x, y)	(0.9623,0.0381)	(0.8641, 0.1118)	(0.8344,0.0907)	(0.8046,0.0727)	(0.7908,0.0650)	(0.7770 ,0.0591)
BL1	ψ	0.2227400E-05	0.2334222E-03	0.9310293E-03	0.1379408E-02	0.1553299E-02	0.1658958E-02
	ω	0.0152618	0.3521649	0.9788872	1.5218878	1.9127837	2.2099272
	(x, y)	(0.0377,0.0381)	(0.0832, 0.0781)	(0.0843, 0.1109)	(0.0728,0.1371)	(0.0642,0.1535)	(0.0586,0.1633)

5.3 A 3-D spherical problem: the Global Electric Circuit

In this section, we test the performance of PHS with polynomial augmentation when solving a fully three-dimensional test problem with irregular boundaries. The selected test problem, known as the Global Electric Circuit (GEC), is from the field of atmospheric electricity with the PDE solving for the divergence of a current (gradient of an electric potential u) in a variable conducting material (the atmosphere) due to charge sources (e.g. electrified thunderstorms). A 3D RBF-FD solver was developed [2] using MQ RBFs augmented with a constant. The electric potential within the atmosphere $u(r, \theta, \lambda)$ is modeled by the 3-D elliptic PDE

$$-\nabla \cdot (\sigma \nabla u) = S, \quad \text{in } \Omega \quad (20)$$

where the variable coefficient $\sigma(r, \theta, \lambda)$ represents the conductivity of the atmosphere and $S(r, \theta, \lambda)$ the source distribution. This equation is obtained by applying Ohm's law to the steady-state current continuity equation. The 3-D domain Ω is defined as $-90^\circ \leq \theta \leq 90^\circ$, $-180^\circ \leq \lambda < 180^\circ$, $k(\theta, \lambda) \leq r \leq r_b$, where the irregular boundary $k(\theta, \lambda)$ is the Earth's topography and r_b is the altitude from sea level on which the top boundary is enforced.

In our present simplified version of the problem, the variable coefficient is assumed to increase exponentially with altitude r as

$$\sigma(r) = \sigma_0 e^{r/c}, \quad (21)$$

where $c = 6$ km and $\sigma_0 = 5 \cdot 10^{-11}$ S/km. To overcome the possible ill-conditioning of the equation due to the highly variable and exponential dependance of (21), we rewrite (20) as in [2],

$$\Delta u + (\nabla \log \sigma) \cdot \nabla u = f. \quad (22)$$

Dirichlet boundary conditions are enforced on the boundary,

$$u = g(\theta, \lambda), \quad \text{on } \partial\Omega \quad (23)$$

and the functions f and g are computed by assuming the exact solution

$$u(r, \theta, \lambda) = [3 Y_9^{-7}(\theta, \lambda) + 2 Y_6^0(\theta, \lambda)] \sin \left(4\pi \frac{r - r_0}{r_b - r_0} \right). \quad (24)$$

Following the same procedure as in [2], equation (22) is solved in the domain which results after the change of variable

$$r(\xi) = A e^{\beta(\xi - \xi_0)} + B, \quad (25)$$

where A and B are constants determined by the conditions

$$r(\xi_0) = r_0 \quad \text{and} \quad r(\xi_b) = r_b. \quad (26)$$

Here, r_0 is the mean radius of the Earth set to 6400 km and $r_b = 6460$ km. Notice that under this change of variable, the Earth is mapped over a sphere of radius ξ_0 and the radial coordinate is exponentially stretched. In [2], the domain is discretized using two different node layouts: scattered nodes based on a repel algorithm between the Earth's surface and ~ 10 km altitude from sea level, and a nested shell model in the rest of the domain. It leads to truly 3-D stencils in the first region, so that the irregularities of the boundary can be properly resolved, and hybrid FD/RBF-FD scheme in the structured nested shell region.




Figure 11: Equatorial view (left) and north pole view (right) of the Earth mapped on a sphere of radius $\xi_0 = 1.15$, with the elevations are scaled up by a factor of 400. The dots are the nodes which discretize the surface of the Earth for the 3-D node distribution $N = 33,389$, which is the bottom boundary of the domain Ω .

Since our present goal is to study the performance of 3-D scattered approximations based on PHS with polynomial augmentation, we have here solved the problem only in the scattered region, i.e. from the surface of the Earth to ~ 12 km from sea level. Therefore, the Dirichlet boundary conditions (23) are enforced on the bottom boundary $k(\theta, \lambda)$ and the top boundary at 6412 km.

To discretize the domain, we follow a procedure similar to the one described in Section 3.2, but adapted to this 3-D geometry. In this case, the domain is filled up using a nested shell model formed by N_r radially aligned spherical shells, where the number of nodes on each shell N_h has varied so that the horizontal internodal distances are kept constant and equal to the radial one. The nodes are then redistributed using a repel algorithm. Those lying outside the domain are discarded. Figure 11 shows an equatorial view (left) and a north pole view (right) of the surface of the Earth mapped on a sphere of radius $\xi_0 = 1.15$. The dots are the nodes which discretize the bottom boundary for a 3-D node distribution with $N = 33,389$ nodes.

To approximate (22), we have used PHS r^7 augmented with all $\binom{4+3}{4} = 35$ polynomial terms up through degree 4 in 3-D. Following the rule of thumb observed in Section 4, where it was noticed that the stencil size n should be about twice the number of polynomial terms included, the stencil size is set to $n = 70$. Figure 12 shows the sparsity pattern after applying sparse reverse Cuthill-McKee ordering (left) and the eigenvalue distribution (right) of the resulting differentiation matrix that approximates (22). All the eigenvalues lie very close to the negative real axis. The linear system of equations is solved as described in Section 3.3, by using a BiCGSTAB iterative method with incomplete LU factorization as preconditioner.

The ℓ_2 error and number of BiCGSTAB iterations are listed in Table 5 as a function of the number of nodes N when solving this PDE problem. These results are plotted in Figure 13. The convergence is $O(h^4)$ and the number of iterations increases linearly as a function of $\sqrt[3]{N}$. We note that 4th degree with $n = 70$ appears to be a good trade-off between accuracy and computational cost. Achieving 6th order convergence would require including up to $\binom{6+3}{6} = 84$ polynomial terms, with a stencil of around $n = 168$ nodes. Hence, the computational cost of computing the differentiation




Figure 12: Sparsity pattern after applying sparse reverse Cuthill-McKee ordering (left) and eigenvalue distribution (right) of the differentiation matrix that approximates (22) using PHS r^7 augmented with polynomials up to 4th degree and $n = 70$, for $N = 33,389$. Note here the different scales on the two axes of the right plot.

Table 5: The specifications for the results shown in Figure 13

N	ℓ_2 error	# iterations
33,389	$1.4 \cdot 10^{-2}$	10
64,048	$8.5 \cdot 10^{-3}$	12
109,455	$4.8 \cdot 10^{-3}$	15.5
255,272	$1.6 \cdot 10^{-3}$	20.5
494,010	$6.4 \cdot 10^{-4}$	25
1,076,812	$2.1 \cdot 10^{-4}$	30.5
1,997,495	$8.6 \cdot 10^{-5}$	38.5

matrices would be about $(168/70)^3 \approx 14$ times higher.

6 Conclusions

A novel RBF-FD method based on Polyharmonic Splines (PHS) with polynomial augmentation is considered in this paper for the numerical solution of elliptic PDE problems. Several test problems are solved to assess the performance of the method. In all of them, various beneficial features for the solution of elliptic PDEs are observed.

In particular, the drawbacks on accuracy and stability due to Runge's phenomenon are overcome by just increasing the stencil size. As a rule of thumb, we have found that, for the numerical method to be stable when dealing with boundaries (without any special treatment such as ghost nodes), the stencil size should be about twice or more the number of polynomial terms. The underlying reason is related with the fact that RBFs with polynomial augmentation is actually a minimization problem, where the RBF-FD weights are the solution of the underdetermined linear

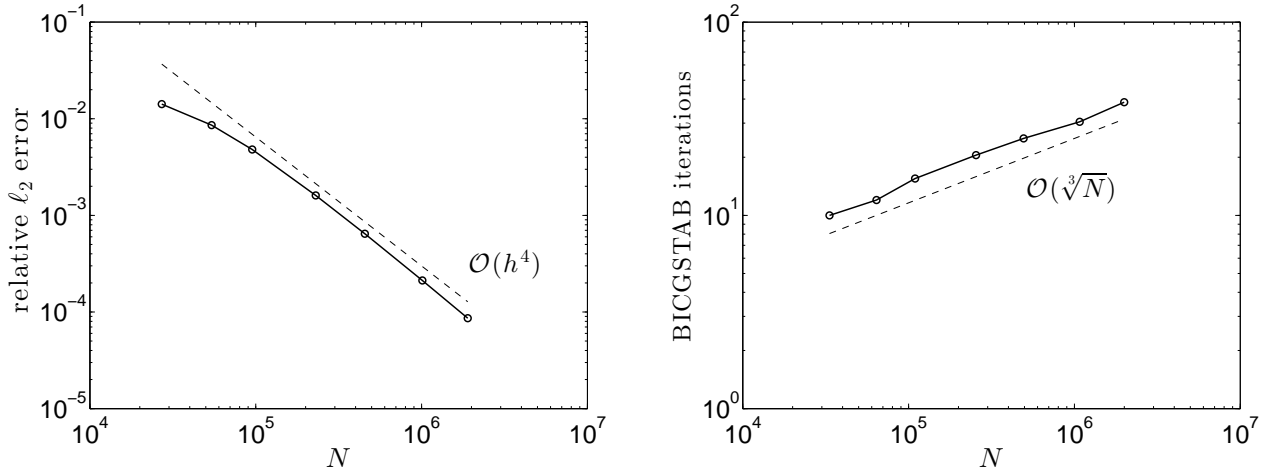


Figure 13: The relative ℓ_2 error (left) and the number of iterations (right) as a function of N for the 3-D global electric circuit test problem.

system $P^T \mathbf{w} = L\mathbf{p}$ minimizing a certain functional $J(\mathbf{w})$. As a result, the convergence order is determined by the polynomial degree, while the stencil size mostly influences the stability of the method (decreasing the stencil weights in magnitude wherever this is possible, given that the order of accuracy needs to be preserved).

In general, this method was found to be very competitive, having all the advantages of RBF-FD methods (high-order convergence, allowing for local node refinement, geometric flexibility, easy to program, etc), while bypassing the main drawbacks, such as the selection of a shape parameter, ill-conditioning and stagnation errors.

Acknowledgments: The National Center for Atmospheric Research is sponsored by the NSF. Most part of this work was done while Victor Bayona was a post-doctoral fellow funded by the Advanced Study Program at the National Center for Atmospheric Research.

References

- [1] E. Barragy and G.F. Carey, *Stream function-vorticity driven cavity solution using p finite elements*, *Comp. & Fluids* **26** (1997), no. 5, 453–468.
- [2] V. Bayona, N. Flyer, G.M. Lucas, and A.J.G. Baumgaertner, *A 3-D RBF-FD solver for modeling the atmospheric global electric circuit with topography (GEC-RBFFD v1.0)*, *Geo. Model Dev.* **8** (2015), 3007–3020.
- [3] O. Botella and R. Peyret, *Benchmark spectral results on the lid-driven cavity flow*, *Comp. Fluids* **27** (1998), no. 4, 421–433.
- [4] P. P. Chinchapatnam, K. Djidjeli, P. B. Nair, and M. Tan, *A compact RBF-FD based meshless method for the incompressible Navier-Stokes equations*, *J. Eng. Maritime Env.* **223** (2009), no. 10, 275–290.
- [5] T.A. Driscoll and B. Fornberg, *Block pseudospectral methods for Maxwell’s equations II: Two-dimensional, discontinuous-coefficient case*, *SIAM J. Sci. Comp.* **21** (1999), no. 3, 1146–1167.

- [6] E. Erturk, T.C. Corke, and C. Gökçöl, *Numerical solutions of 2-D steady incompressible driven cavity flow at high Reynolds numbers*, Int. J. Num. Meth. Fluids **48** (2005), no. 7, 747–774.
- [7] G. E. Fasshauer, *Meshfree Approximation Methods with MATLAB*, Interdisciplinary Mathematical Sciences - Vol. 6, World Scientific Publishers, Singapore, 2007.
- [8] A. I. Fedoseyev, M. J. Friedman, and E. J. Kansa, *Improved multiquadric method for elliptic partial differential equations via PDE collocation on the boundary*, Comput. Math. Appl. **43** (2002), 439–455.
- [9] N. Flyer, G. A. Barnett, and L. J. Wicker, *Enhancing finite differences with radial basis functions: Experiments on the Navier–Stokes equations*, J. Comput. Phys. **39–62** (2016), 39–62.
- [10] N. Flyer, B. Fornberg, V. Bayona, and G. Barnett, *On the role of polynomials in RBF-FD approximations: I. Interpolation and accuracy*, J. Comput. Phys. **321** (2016), 21–38.
- [11] N. Flyer, G. B. Wright, and B. Fornberg, *Radial basis function-generated finite differences: A mesh-free method for computational geosciences*, Handbook of Geomathematics (Berlin) (W. Freeden, M.Z. Nashed, and T. Sonar, eds.), Springer Verlag, 2014.
- [12] B. Fornberg, *A pseudospectral fictitious point method for high order initial-boundary value problems*, SIAM J. on Sci.Comput. **28** (2006), no. 5, 1716–1729.
- [13] B. Fornberg and N. Flyer, *Fast generation of 2-D node distributions for mesh-free PDE discretizations*, Comp. Math. Appl. **69** (2015), no. 7, 531–544.
- [14] B. Fornberg and N. Flyer, *A Primer on Radial Basis Functions with Applications to the Geosciences*, CBMS-NSF Regional Conf. Series in Applied Math, vol. 87, SIAM Press, 2015.
- [15] ———, *Solving PDEs with radial basis functions*, Acta Numerica **25** (2015), 215–258.
- [16] U.K.N.G. Ghia, K.N. Ghia, and C.T. Shin, *High-Re solutions for incompressible flow using the Navier-Stokes equations and a multigrid method*, J. Comput. Phys. **48** (1982), no. 3, 387–411.
- [17] A. Iske, *On the approximation order and numerical stability of local Lagrange interpolation by polyharmonic splines*, Modern Developments in Multivariate Approximation (W. Haussmann, K. Jetter, M. Reimer, and J. Stöckler, eds.), International Series of Numerical Mathematics, Birkhäuser Verlag, 2003, pp. 153–165.
- [18] E. Larsson and B. Fornberg, *A numerical study of some radial basis function based solution methods for elliptic PDEs*, Comput. Math. Appl. **46** (2003), 891–902.
- [19] M. Li, T. Tang, and B. Fornberg, *A compact fourth-order finite difference scheme for the steady incompressible Navier-Stokes equations*, Int. J. Num. Meth. Fluids **20** (1995), no. 10, 1137–1151.
- [20] P.O. Persson and G. Strang, *A simple mesh generator in MATLAB*, SIAM Review **46** (2004), no. 2, 329–345.
- [21] Y.V.S.S. Sanyasiraju and G. Chandhini, *Local radial basis function based gridfree scheme for unsteady incompressible viscous flows*, J.Comput. Phys. **227** (2008), no. 20, 8922–8948.
- [22] G.D. Smith, *Numerical Solution of Partial Differential Equations*, Oxford mathematical handbooks, Oxford University Press, 1965.