

1-1-2017

A Radial Basis Function (RBF) Compact Finite Difference (FD) Scheme for Reaction-Diffusion Equations on Surfaces

Erik Lehto

KTH Royal Institute of Technology

Varun Shankar

University of Utah

Grady B. Wright

Boise State University

A RADIAL BASIS FUNCTION (RBF) COMPACT FINITE DIFFERENCE (FD) SCHEME FOR REACTION-DIFFUSION EQUATIONS ON SURFACES*

ERIK LEHTO[†], VARUN SHANKAR[‡], AND GRADY B. WRIGHT[§]

Abstract. We present a new high-order, local meshfree method for numerically solving reaction diffusion equations on smooth surfaces of codimension 1 embedded in \mathbb{R}^d . The novelty of the method is in the approximation of the Laplace–Beltrami operator for a given surface using Hermite radial basis function (RBF) interpolation over local node sets on the surface. This leads to compact (or implicit) RBF generated finite difference (RBF-FD) formulas for the Laplace–Beltrami operator, which gives rise to sparse differentiation matrices. The method only requires a set of (scattered) nodes on the surface and an approximation to the surface normal vectors at these nodes. Additionally, the method is based on Cartesian coordinates and thus does not suffer from any coordinate singularities. We also present an algorithm for selecting the nodes used to construct the compact RBF-FD formulas that can guarantee the resulting differentiation matrices have desirable stability properties. The improved accuracy and computational cost that can be achieved with this method over the standard (explicit) RBF-FD method are demonstrated with a series of numerical examples. We also illustrate the flexibility and general applicability of the method by solving two different reaction-diffusion equations on surfaces that are defined implicitly and only by point clouds.

Key words. RBF-FD, RBF-HFD, manifolds, reaction diffusion

AMS subject classifications. 41A21, 65D05, 65E05, 65F22

DOI. 10.1137/16M1095457

1. Introduction. Global radial basis function (RBF) methods are quite popular for the numerical solution of various partial differential equations (PDEs) due to their ability to handle scattered node layouts, their simplicity of implementation, and their potential for spectral accuracy for smooth problems. These methods have been successfully applied to the solution of PDEs in various geometries in \mathbb{R}^2 and \mathbb{R}^3 (e.g., [12, 17]), including spherical domains (e.g., [27, 16, 42]), and more general surfaces embedded in \mathbb{R}^3 (e.g., [26, 36]).

When high orders of algebraic accuracy are sufficient for a given problem, or if the solutions to the problem are expected to only have finite smoothness, RBF generated finite difference (RBF-FD) formulas are an attractive alternative to global RBFs as they perform better in terms of accuracy per computational cost [17]. These formulas are generated from RBF interpolation over *local* sets of nodes (stencils) so that the resulting differentiation matrices are sparse as in the standard FD method. In contrast to standard FD methods, however, the RBF-FD method can naturally handle irregular geometries and scattered node layouts. Additionally, their locality

*Submitted to the journal’s Methods and Algorithms for Scientific Computing section September 23, 2016; accepted for publication (in revised form) March 27, 2017; published electronically September 21, 2017.

<http://www.siam.org/journals/sisc/39-5/M109545.html>

Funding: The work of the second author was supported by this project under NSF-DMS 1160432. The work of the third author was funded by support for this project under grants NSF-DMS 1160379 and NSF-ACI 1440638.

[†]Corresponding author. Department of Mathematics, KTH Royal Institute of Technology, Stockholm, 10044, Sweden (elehto@kth.se).

[‡]Department of Mathematics, University of Utah, Salt Lake City, UT 84112 (vshankar@math.utah.edu).

[§]Department of Mathematics, Boise State University, Boise, ID 837235 (gradywright@boisestate.edu).

makes them more flexibility in terms of local refinement strategies than global RBF methods. The strength of the RBF-FD method has been leveraged to solve problems on planar domains, e.g., [39, 6, 43, 7, 40], the surface of a sphere [19, 15], and, more recently, very general surfaces represented solely by point clouds and normal vectors [38].

It is natural to view these two classes of RBF methods as extensions of classical methods to scattered nodes and irregular geometries. The global RBF method for surface PDEs in [26] may be viewed as an extension of polynomial based (or Fourier based) pseudospectral methods to surfaces, while the RBF-FD method presented in [38] may be viewed as an extension of standard, polynomial based FD methods to surfaces. In this work, we turn our attention to the extension of a third important class of classical methods to surfaces: the so-called compact, implicit, or Hermite FD methods, first introduced by Collatz [9]. We use the acronym HFD for these schemes to avoid the obvious confusion with CFD, and because they will ultimately be based on Hermite interpolation.

The goal of HFD methods is to solve a given PDE numerically by computing more accurate approximations to the differential operators in the PDE. In these schemes, this improved accuracy is obtained by using additional information from the PDE itself, rather than increasing the stencil size, as is the usual way to increase the accuracy with standard (or explicit) FD methods. HFD schemes are thus typically more computationally efficient than standard FD schemes, as they can obtain higher accuracy and resolution for the same stencil size [31]. Further, the differentiation matrices obtained from HFD formulas often have desirable properties such as diagonal dominance, leading to both enhanced numerical stability and faster convergence of iterative methods used in solving the sparse linear systems that arise when using these matrices to discretize a PDE. While HFD schemes have already been successfully generalized to scattered node layouts [43], the application of these schemes to the solution of PDEs on surfaces presents significant challenges due to the presence of surface differential operators. In this article, we overcome those challenges and present a new RBF-HFD scheme for the solution of reaction-diffusion equations on surfaces.¹ The resulting method uses Cartesian coordinates, thereby avoiding the singularities typically associated with intrinsic coordinate systems. Further, our new method only uses nodes on the surface in consideration, making it more computationally efficient than embedded narrow-band methods that solve the PDE in a narrow band in the embedding space (e.g., [32, 36]). Finally, the RBF-HFD formulas require fewer nodes than the RBF-FD method presented in [38] for the same accuracy, while also possessing improved stability properties.

The remainder of the paper is organized as follows. In section 2, we briefly review the formulation of surface differential operators in Cartesian coordinates. Section 3 outlines Hermite RBF interpolation on scattered node sets in \mathbb{R}^d . Section 4 describes how to use approximations to the Hermite RBF interpolants to generate RBF-HFD weights for approximating the surface Laplacian, and also how these can be arranged into *sparse* differentiation matrices. We follow this in section 5 with a brief discussion of how to use these differentiation matrices in a method-of-lines formulation for numerically solving forced diffusion equations on surfaces. In section 6, we discuss the stability of our method by studying the Gershgorin sets associated with the eigenvalues of our differentiation matrices. In section 7, we numerically demonstrate the

¹Throughout this paper, we will use the terms surface or manifold to refer to embedded submanifolds of codimension 1 in \mathbb{R}^d with no boundary, and focus on the specific case of $d = 3$.

accuracy and efficiency of our method for the forced scalar diffusion equation on two different surfaces. We also present a few applications of our method to two species reaction diffusion equations on implicitly defined surfaces and surfaces defined by point clouds, which have relevant biological applications. We conclude our paper with a summary and discussion of future research directions in section 8.

2. Review of differential geometry. While the standard way of expressing differential operators on surfaces is through the use of intrinsic coordinates, covariant derivatives and metric tensors [5], we instead choose to formulate these operators entirely in Cartesian (or extrinsic) coordinates, as this avoids any singularities associated with intrinsic coordinate systems. Consider the standard gradient operator in \mathbb{R}^3 , $\nabla = [\partial_x \ \partial_y \ \partial_z]^T$. If we apply this to a differentiable function f at a point $\mathbf{x} = (x, y, z)$ on the surface \mathbb{M} and then project the resulting vector into the tangent space of the surface, then this gives the surface gradient of f , which we denote as $\nabla_{\mathbb{M}}f$. Mathematically, this can be accomplished as follows. Let $\mathbf{n} = [n^x \ n^y \ n^z]^T$ be the *unit* normal vector to \mathbb{M} at \mathbf{x} ; then

$$\nabla_{\mathbb{M}}f = \nabla f - \mathbf{n}(\mathbf{n} \cdot \nabla f) = \nabla f - \mathbf{nn}^T(\nabla f).$$

Thus, the surface gradient operator can be written entirely in Cartesian coordinates as

$$\nabla_{\mathbb{M}} := \nabla - \mathbf{nn}^T \nabla = \underbrace{(\mathcal{I} - \mathbf{nn}^T)}_{\mathcal{P}} \nabla,$$

where \mathcal{I} is the 3-by-3 identity matrix. \mathcal{P} is a *projection operator* that takes a vector field in \mathbb{R}^3 sampled at a point \mathbf{x} on the surface and projects it onto the tangent plane to the surface at \mathbf{x} . An explicit expression for this operator is given by

$$(1) \quad \mathcal{P} = \begin{bmatrix} (1 - n^x n^x) & -n^x n^y & -n^x n^z \\ -n^x n^y & (1 - n^y n^y) & -n^y n^z \\ -n^x n^z & -n^y n^z & (1 - n^z n^z) \end{bmatrix} = [\mathbf{p}^x \ \mathbf{p}^y \ \mathbf{p}^z],$$

where \mathbf{p}^x , \mathbf{p}^y , and \mathbf{p}^z are vectors representing the projection operators in the x , y , and z directions, respectively. We can now use \mathbf{p}^x , \mathbf{p}^y , and \mathbf{p}^z to obtain the following (more convenient) expression for $\nabla_{\mathbb{M}}$:

$$(2) \quad \nabla_{\mathbb{M}} := \mathcal{P}\nabla = \begin{bmatrix} \mathbf{p}^x \cdot \nabla \\ \mathbf{p}^y \cdot \nabla \\ \mathbf{p}^z \cdot \nabla \end{bmatrix} = \begin{bmatrix} \mathcal{G}^x \\ \mathcal{G}^y \\ \mathcal{G}^z \end{bmatrix},$$

where \mathcal{G}^x , \mathcal{G}^y , and \mathcal{G}^z are the components of the surface gradient along each of the coordinate directions in \mathbb{R}^3 . Now, the surface Laplace (Laplace–Beltrami) operator $\Delta_{\mathbb{M}}$ can be obtained by applying the surface divergence to the surface gradient [26]. This can naturally be expressed using \mathcal{G}^x , \mathcal{G}^y , and \mathcal{G}^z as

$$(3) \quad \Delta_{\mathbb{M}} := \nabla_{\mathbb{M}} \cdot \nabla_{\mathbb{M}} = (\mathcal{P}\nabla) \cdot (\mathcal{P}\nabla) = \mathcal{G}^x \mathcal{G}^x + \mathcal{G}^y \mathcal{G}^y + \mathcal{G}^z \mathcal{G}^z.$$

This gives an explicit expression for the surface Laplace operator entirely in Cartesian components. We will use this expression in our numerical approximation to the surface Laplacian.

3. Hermite Interpolation with RBFs. We now review Hermite interpolation with RBFs, a technique essential to deriving the new RBF-HFD scheme outlined in the next section. Let $\Omega \subseteq \mathbb{R}^d$ and let $\phi : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$ be a scalar-valued radial kernel, i.e., $\phi(\mathbf{x}, \mathbf{y}) := \phi(\|\mathbf{x} - \mathbf{y}\|)$ for $\mathbf{x}, \mathbf{y} \in \Omega$, where $\|\cdot\|$ is the standard Euclidean norm in \mathbb{R}^d . Let \mathcal{L} be a linear functional and suppose we are given samples of a continuous target function f at a set of distinct nodes $X = \{\mathbf{x}_i\}_{i=1}^n \subset \Omega$ and samples of $\mathcal{L}f$ at a set of distinct nodes $\tilde{X} = \{\tilde{\mathbf{x}}_j\}_{j=1}^m \subset \Omega$. Then we consider the following Hermite RBF interpolant to the this data, proposed first by Wu [44]:

$$(4) \quad I_\phi f(\mathbf{x}) = \sum_{i=1}^n c_i \phi(\mathbf{x}, \mathbf{x}_i) + \sum_{j=1}^m d_j \mathcal{L}_2 \phi(\mathbf{x}, \tilde{\mathbf{x}}_j) + \alpha.$$

Here we have used the notation \mathcal{L}_2 to mean that \mathcal{L} is applied to ϕ with respect to its second argument. Later we will similarly use \mathcal{L}_1 to mean that \mathcal{L} is applied to ϕ with respect to its first argument. The expansion coefficients $\{c_i\}_{i=1}^n$ and $\{d_j\}_{j=1}^m$ in (4) are determined by enforcing the (Hermite) interpolation conditions

$$(5) \quad I_\phi f|_X = f|_X,$$

$$(6) \quad \mathcal{L}(I_\phi f)|_{\tilde{X}} = (\mathcal{L}f)|_{\tilde{X}},$$

while the constant α is obtained by enforcing the moment condition $\sum_{i=1}^n c_i = 0$. These conditions can be represented as the following block linear system:

$$(7) \quad \underbrace{\begin{bmatrix} A & B_2 & \mathbf{e} \\ B_1 & C & \mathbf{0} \\ \mathbf{e}^T & \mathbf{0}^T & 0 \end{bmatrix}}_{A_H} \begin{bmatrix} \mathbf{c} \\ \mathbf{d} \\ \alpha \end{bmatrix} = \begin{bmatrix} \mathbf{f} \\ \mathcal{L}\mathbf{f} \\ 0 \end{bmatrix},$$

where

$$\begin{aligned} A_{i,j} &= \phi(\mathbf{x}_i, \mathbf{x}_j), & i, j &= 1, \dots, n, \\ (B_2)_{i,j} &= \mathcal{L}_2 \phi(\mathbf{x}_i, \tilde{\mathbf{x}}_j), & i &= 1, \dots, n, \quad j = 1, \dots, m, \\ (B_1)_{i,j} &= \mathcal{L}_1 \phi(\tilde{\mathbf{x}}_i, \mathbf{x}_j), & i &= 1, \dots, m, \quad j = 1, \dots, n, \\ C_{i,j} &= \mathcal{L}_1 \mathcal{L}_2 \phi(\tilde{\mathbf{x}}_i, \tilde{\mathbf{x}}_j), & i, j &= 1, \dots, m, \\ \mathbf{e}_i &= 1, & i &= 1, \dots, n. \end{aligned}$$

Because ϕ is radially symmetric, we have that $\phi(\mathbf{x}, \tilde{\mathbf{x}}) = \phi(\tilde{\mathbf{x}}, \mathbf{x})$, so that $\mathcal{L}_2 \phi(\mathbf{x}, \tilde{\mathbf{x}}) = \mathcal{L}_1 \phi(\tilde{\mathbf{x}}, \mathbf{x})$. This means that $A = A^T$, $C = C^T$, and $B_2 = B_1^T$ so that the matrix A_H is symmetric. If ϕ is, for example, positive definite or order-1 conditionally positive definite, then under very mild conditions on \mathcal{L} the linear system (7) is nonsingular [44, 33].

We will also make use of regular RBF interpolation, which consists only of interpolating function values, in the subsequent section. These interpolants are simply given by (4) with m set equal to zero and the constant α omitted, i.e.,

$$(8) \quad I_\phi f(\mathbf{x}) = \sum_{i=1}^n c_i \phi(\mathbf{x}, \mathbf{x}_i).$$

In this case, we only enforce the conditions (5), which can be represented by the linear system

$$(9) \quad A_R \mathbf{c} = \mathbf{f},$$

where A_R is the same matrix as A in (7). We use the subscript R to denote the linear system for the *regular* interpolant as opposed to the subscript H in (7) for the linear system associated with the *Hermite* interpolant.

In this study, the interpolation nodes X and “functional nodes” \tilde{X} lie on an embedded lower dimensional surface $\Omega = \mathbb{M}$ in \mathbb{R}^d . However, we will still use the standard Euclidean distance in \mathbb{R}^d for computing $\phi(\mathbf{x}, \tilde{\mathbf{x}}) = \phi(\|\mathbf{x} - \tilde{\mathbf{x}}\|)$ in (4) (i.e., straight line distances rather than distances intrinsic to the surface). A theoretical foundation for RBF interpolation on surfaces with the straight-line distance measure is given in [24], along with proofs of favorable error estimates.

While it is possible to use any (conditionally) positive-definite kernel within the RBF-FD and RBF-HFD methods (e.g., [39, 43, 4, 10, 15]), we use the Gaussian (GA) kernel, which is positive definite in \mathbb{R}^d for any d . All infinitely smooth kernels feature a shape parameter ε such that large values of ε make the kernels peaked, while smaller ε values make them flat. In the limit as $\varepsilon \rightarrow 0$, Gaussian RBF interpolants to data converge to (multivariate) polynomial interpolants in \mathbb{R}^d [11, 29, 37], and to spherical harmonic interpolants on the sphere \mathbb{S}^2 [21]. While smaller values of ε generally lead to greater accuracy for smooth target functions [22, 29], the interpolation matrix in (7) becomes increasingly ill-conditioned as $\varepsilon \rightarrow 0$ (see, e.g., [23]). Some stable algorithms have been developed for bypassing this ill-conditioning [22, 21, 13, 18, 20], but these algorithms typically break down when the data sites lie on a submanifold $\mathbb{M} \subset \mathbb{R}^d$, as in the present study, due to nodes being nonunisolvent with respect to polynomials in \mathbb{R}^d . While some strategies have recently been undertaken to resolve them in \mathbb{R}^d [30], a robust approach is not yet available for surfaces.

4. RBF-HFD formulas for the surface Laplacian. Let $\Xi = \{\xi_k\}_{k=1}^N$ denote a set of (scattered) node locations on a surface \mathbb{M} of dimension 2 embedded in \mathbb{R}^3 and suppose $f : \mathbb{M} \rightarrow \mathbb{R}$ is some differentiable function sampled on Ξ . Our goal is to approximate $\Delta_{\mathbb{M}}f|_{\Xi}$ with HFD-style local approximations to the operator $\Delta_{\mathbb{M}}$. Without loss of generality, let the node where we want to approximate $\Delta_{\mathbb{M}}f$ at be ξ_1 , and let ξ_2, \dots, ξ_p be the $p - 1$ nearest neighbors to ξ_1 , measured by Euclidean distance in \mathbb{R}^3 . We refer to ξ_1 and its $p - 1$ nearest neighbors as the neighborhood of ξ_1 on the surface and denote this neighborhood as $S_1 = \{\xi_\ell\}_{\ell=1}^p$; this neighborhood will comprise the candidate nodes that make up the HFD stencil for ξ_1 . We seek an approximation to $\Delta_{\mathbb{M}}f$ at ξ_1 that involves a linear combination of the values of f and $\Delta_{\mathbb{M}}f$ over some subset of nodes from S_1 of the form

$$(10) \quad (\Delta_{\mathbb{M}}f)|_{\mathbf{x}=\xi_1} \approx \sum_{i \in J} w_i f(\xi_i) + \sum_{j \in \tilde{J}} \tilde{w}_j (\Delta_{\mathbb{M}}f)|_{\mathbf{x}=\xi_j},$$

where J and \tilde{J} denote index sets of size $n \leq p$ and $m < p$, respectively, into S_1 for the explicit and the implicit (or Hermite) part of the stencil, respectively. We will assume that $1 \in J$, but $1 \notin \tilde{J}$ (otherwise a trivial solution would exist). Using the notation of the previous section, we will let the n nodes indicated by J be denoted by $X = \{\mathbf{x}_i\}_{i=1}^n$ and the m nodes indicated by \tilde{J} be denoted by $\tilde{X} = \{\tilde{\mathbf{x}}_j\}_{j=1}^m$. Additionally, we always set $\mathbf{x}_1 = \xi_1$. Using this notation we can rewrite (10) as

$$(11) \quad (\Delta_{\mathbb{M}}f)|_{\mathbf{x}=\mathbf{x}_1} \approx \sum_{i=1}^n w_i f(\mathbf{x}_i) + \sum_{j=1}^m \tilde{w}_j (\Delta_{\mathbb{M}}f)|_{\mathbf{x}=\tilde{\mathbf{x}}_j}.$$

The weights $\{w_i\}_{i=1}^n$ and $\{\tilde{w}_j\}_{j=1}^m$ in this approximation will be computed using RBFs, and will be referred to as RBF-HFD weights.

4.1. Computation of the weights from the Hermite interpolant. The method from [43] determines the RBF-HFD weights in (11) from the Hermite RBF interpolant (4) constructed with $\mathcal{L} = \Delta_{\mathbb{M}}$. To compute the weights, consider the problem of applying $\Delta_{\mathbb{M}}$ to the interpolant (4) and evaluating it at \mathbf{x}_1 to approximate $\Delta_{\mathbb{M}}f|_{\mathbf{x}=\mathbf{x}_1}$. The resulting approximation would be exact whenever f is any of the functions $\phi(\mathbf{x}, \mathbf{x}_i), i = 1, \dots, n, \mathcal{L}_2\phi(\mathbf{x}, \tilde{\mathbf{x}}_j) = \Delta_{\mathbb{M},2}\phi(\mathbf{x}, \tilde{\mathbf{x}}_j), j = 1, \dots, m$, or a nonzero constant (since the interpolant is exact for these f). Thus, the weights $\{w_i\}_{i=1}^n$ and $\{\tilde{w}_j\}_{j=1}^m$ are the values that make (11) exact for these values of f . This can be written as the following linear system:

$$(12) \quad \underbrace{\begin{bmatrix} A & B & \mathbf{e} \\ B^T & C & \mathbf{0} \\ \mathbf{e}^T & \mathbf{0}^T & 0 \end{bmatrix}}_{A_H} \begin{bmatrix} w \\ \tilde{w} \\ \alpha \end{bmatrix} = \begin{bmatrix} \Delta_{\mathbb{M},1}\phi(\mathbf{x}, \mathbf{x}_i)|_{\mathbf{x}=\mathbf{x}_1} \\ \Delta_{\mathbb{M},1}\Delta_{\mathbb{M},2}\phi(\mathbf{x}, \tilde{\mathbf{x}}_j)|_{\mathbf{x}=\mathbf{x}_1} \\ 0 \end{bmatrix},$$

where the block matrices A and C are the same as those given in the Hermite interpolation matrix (7), $B = B_2 = B_1^T$ in this same matrix (recall that the matrix in (7) is symmetric), $\Delta_{\mathbb{M},1} = \mathcal{L}_1$ and $\Delta_{\mathbb{M},2} = \mathcal{L}_2$, and $i = 1, \dots, n$ and $j = 1, \dots, m$ to form block vectors of length n and m in the right-hand side. Note that the constant α is not used for anything in the actual RBF-HFD formula.

The issue with using (12) for determining the RBF-FD weights is that one has to explicitly compute $\Delta_{\mathbb{M},1}\phi(\mathbf{x}, \tilde{\mathbf{x}}_j)$ and $\Delta_{\mathbb{M},1}\Delta_{\mathbb{M},2}\phi(\mathbf{x}, \tilde{\mathbf{x}}_j)$. As discussed in section 2, constructing $\Delta_{\mathbb{M}}$ requires having explicit information about the underlying surface, such as an analytical expression for the surface normal vectors. Even in cases where these are known, the resulting formulas for computing $\Delta_{\mathbb{M},1}\phi$ and $\Delta_{\mathbb{M},1}\Delta_{\mathbb{M},2}\phi$ are likely to be quite complex. Moreover, we are interested in surfaces that are defined by point clouds and where only numerical representations of the normal vectors are available. Thus, constructing the system (12) analytically will not be possible. However, it is possible to construct an approximation to the entries of this system using the regular RBF-FD method from [38], which is based on iterated differentiation (see also [25]). This is the approach we take.

4.2. Computation of the weights from iterated differentiation. The first goal is to compute approximations of the entries in B^T in (12) and the entries of the first vector block in the right-hand side of this equation. We state the entries of B^T explicitly as it will help elucidate the discussion of their approximation:

$$(13) \quad B^T = \begin{bmatrix} \Delta_{\mathbb{M},1}\phi(\tilde{\mathbf{x}}_1, \mathbf{x}_1) & \cdots & \Delta_{\mathbb{M},1}\phi(\tilde{\mathbf{x}}_1, \mathbf{x}_n) \\ \vdots & \ddots & \vdots \\ \Delta_{\mathbb{M},1}\phi(\tilde{\mathbf{x}}_m, \mathbf{x}_1) & \cdots & \Delta_{\mathbb{M},1}\phi(\tilde{\mathbf{x}}_m, \mathbf{x}_n) \end{bmatrix}.$$

We compute these approximations by constructing an approximation to $\Delta_{\mathbb{M}}$ using discrete approximations to $\mathcal{G}^x, \mathcal{G}^y$, and \mathcal{G}^z in (2) computed from the standard RBF interpolant (8) over the candidate stencil nodes $S_1 = \{\boldsymbol{\xi}_k\}_{k=1}^p$. To this end, consider, for example, applying \mathcal{G}^x to the interpolant $I_\phi f$ in (8) based on the nodes in S_1 (here the target function f is not important) and then evaluating it at S_1 :

$$(14) \quad (\mathcal{G}^x I_\phi f(\mathbf{x}))|_{\mathbf{x}=\boldsymbol{\xi}_i} = \sum_{j=1}^p c_j \underbrace{(\mathcal{G}^x \phi(\mathbf{x}, \boldsymbol{\xi}_j))|_{\mathbf{x}=\boldsymbol{\xi}_i}}_{D_{ij}^x}, \quad i = 1, \dots, p.$$

We can rewrite (14) so that it explicitly depends only on the vector of samples $f|_{S_1}$ using (9) as follows:

$$(15) \quad (\mathcal{G}^x I_\phi f)|_{S_1} = D^x c_f = D^x A_R^{-1} f|_{S_1} =: G^x f|_{S_1}.$$

Here G^x is a p -by- p differentiation matrix that represents the RBF approximation to the x component of the surface gradient operator over the set of nodes in S_1 . Now, letting

$$(16) \quad D_{i,j}^y = (\mathcal{G}^y \phi(\mathbf{x}, \boldsymbol{\xi}_j))|_{\mathbf{x}=\boldsymbol{\xi}_i} \quad \text{and} \quad D_{i,j}^z = (\mathcal{G}^z \phi(\mathbf{x}, \boldsymbol{\xi}_j))|_{\mathbf{x}=\boldsymbol{\xi}_i}, \quad i, j = 1, \dots, p,$$

we can obtain similar approximations to the y and z components of the surface gradient operator on S_1 as

$$(17) \quad (\mathcal{G}^y I_\phi f)|_{S_1} = D^y A_R^{-1} f|_{S_1} =: G^y f|_{S_1},$$

$$(18) \quad (\mathcal{G}^z I_\phi f)|_{S_1} = D^z A_R^{-1} f|_{S_1} =: G^z f|_{S_1}.$$

To obtain an approximation to Δ_M at the candidate stencil nodes S_1 , we mimic the continuous formulation of the surface Laplacian in (3), replacing the continuous operators \mathcal{G}^x , \mathcal{G}^y , and \mathcal{G}^z with the differentiation matrices G^x , G^y , and G^z , respectively. This gives the following differentiation matrix for approximating the surface Laplacian at the nodes S_1 :

$$(19) \quad \begin{aligned} L_{M,1} &= G^x G^x + G^y G^y + G^z G^z \\ &= \underbrace{(D^x A_R^{-1} D^x + D^y A_R^{-1} D^y + D^z A_R^{-1} D^z)}_{\hat{B}^T} A_R^{-1}. \end{aligned}$$

When applying $L_{M,1}$ to a vector of samples of a target function f taken over S_1 , this is equivalent to interpolating the target function with the regular RBF interpolant (8), computing the components of the surface gradient of the interpolant, then interpolating each of these components again using (8), applying the surface divergence, then evaluating this at the nodes in S_1 . This is a type of iterated derivative approximation [25] and has the advantage of not needing the explicit formulas for the normal vectors (or their derivatives) to the surface M .

Recall that the node sets X and \tilde{X} are subsets of S_1 given by the index sets J and \tilde{J} , respectively (cf. (10)). Thus, to approximate the (i, ℓ) entry, $\Delta_{M,1} \phi(\tilde{\mathbf{x}}_i, \mathbf{x}_\ell)$, of B^T in (13), we can first apply $L_{M,1}$ to the vector of samples of $\phi(\mathbf{x}, \mathbf{x}_\ell)$ at S_1 ,

$$(20) \quad [\phi(\boldsymbol{\xi}_1, \mathbf{x}_\ell) \quad \phi(\boldsymbol{\xi}_2, \mathbf{x}_\ell) \quad \cdots \quad \phi(\boldsymbol{\xi}_p, \mathbf{x}_\ell)]^T,$$

which gives a vector containing approximations to $\Delta_{M,1} \phi(\boldsymbol{\xi}_k, \mathbf{x}_\ell)$, $k = 1, \dots, p$. The approximation to $\Delta_{M,1} \phi(\tilde{\mathbf{x}}_i, \mathbf{x}_\ell)$ is then given by the row in this vector corresponding to the i th value in \tilde{J} (which we denote by \tilde{J}_i). Note, however, that the vector (20) is just the J_ℓ column of A_R in (15), so that the approximation to $\Delta_{M,1} \phi(\tilde{\mathbf{x}}_i, \mathbf{x}_\ell)$ obtained from applying $L_{M,1}$ to (20) is just given by the \tilde{J}_i and J_ℓ column of \hat{B}^T in (19). Thus, all entries in B^T in (13) can be similarly obtained directly from the rows and columns or \hat{B}^T using the index sets J and \tilde{J} . Additionally, the vector in the first block of the right hand side of (12) can be approximated from \hat{B}^T ; in this case, from the first row of \hat{B}^T and from the columns corresponding to J_i , $i = 1, \dots, n$.

The second goal is to compute approximations to the entries of C in (12) and the entries of the second vector block in the right-hand side of this equation. We give the entries of C explicitly to again elucidate the discussion:

$$(21) \quad C = \begin{bmatrix} \Delta_{\mathbb{M},1} \Delta_{\mathbb{M},2} \phi(\tilde{\mathbf{x}}_1, \tilde{\mathbf{x}}_1) & \cdots & \Delta_{\mathbb{M},1} \Delta_{\mathbb{M},2} \phi(\tilde{\mathbf{x}}_1, \tilde{\mathbf{x}}_m) \\ \vdots & \ddots & \vdots \\ \Delta_{\mathbb{M},1} \Delta_{\mathbb{M},2} \phi(\tilde{\mathbf{x}}_m, \tilde{\mathbf{x}}_1) & \cdots & \Delta_{\mathbb{M},1} \Delta_{\mathbb{M},2} \phi(\tilde{\mathbf{x}}_m, \tilde{\mathbf{x}}_m) \end{bmatrix}.$$

To approximate the operator $\Delta_{\mathbb{M},1} \Delta_{\mathbb{M},2}$ we again use iterated differentiation involving the differentiation matrices G^x , G^y , and G^z . Using the idea of (19), we can approximate $\Delta_{\mathbb{M},2}$ at the candidate stencil nodes S_1 using the differentiation matrix

$$(22) \quad L_{\mathbb{M},2} = \underbrace{((D^x)^T A_{\mathbb{R}}^{-1} (D^x)^T + (D^y)^T A_{\mathbb{R}}^{-1} (D^y)^T + (D^z)^T A_{\mathbb{R}}^{-1} (D^z)^T)}_{\hat{B}} A_{\mathbb{R}}^{-1},$$

where D^x , D^y , and D^z are given by (14) and (16). We then approximate $\Delta_{\mathbb{M},1} \Delta_{\mathbb{M},2}$ at the nodes in S_1 as

$$(23) \quad L_{\mathbb{M},1} L_{\mathbb{M},2} = \underbrace{(\hat{B}^T A_{\mathbb{R}}^{-1} \hat{B})}_{\hat{C}} A_{\mathbb{R}}^{-1}.$$

Using similar arguments as above for extracting approximations to the elements of B^T from \hat{B}^T , we can extract approximations to the elements of C from \hat{C} . For example, entry $C_{i,\ell}$ can be approximated by the entry in the \tilde{J}_i row and \tilde{J}_ℓ column of \hat{C} . The elements in the second block vector of the right-hand side of (12) can similarly be extracted from \hat{C} . In practice, \hat{B} and \hat{C} are formed by solving linear systems using the Cholesky factorization of $A_{\mathbb{R}}$, instead of computing $A_{\mathbb{R}}^{-1}$. Note that this ensures that \hat{C} is symmetric so that the approximation to C will also be symmetric.

Upon obtaining approximations to the entries of B^T and C and the vector in the right-hand side of (12), we substitute these into the system (12) and solve it to obtain iterated RBF-HFD weights $\{w\}_{i=1}^n$ and $\{\tilde{w}\}_{j=1}^m$ to be used in (11).

For each node $\xi_k \in \Xi$, $k = 1, \dots, N$, we repeat the above procedure of finding its $p - 1$ nearest neighbors (candidate stencil nodes S_k), selecting index sets for the explicit and implicit stencils, computing approximations to the p -by- p submatrices \hat{B}^T and \hat{C} , and extracting the entries from these matrices to use for solving for the weights in (12). These weights are then arranged into two *sparse* N -by- N matrices L_{Ξ} and \tilde{L}_{Ξ} for approximating the surface Laplacian over all the nodes in Ξ (see section 5 for how L_{Ξ} and \tilde{L}_{Ξ} are used for solving a PDE). Each row of L_{Ξ} has n nonzero entries and each row of \tilde{L}_{Ξ} has m nonzero entries.

The computational cost of computing the weights for node ξ_k is $\mathcal{O}(p^3)$, and there are N such stencils, so that the total cost of computing the entries of L_{Ξ} and \tilde{L}_{Ξ} is $\mathcal{O}(p^3 N)$. In our application of the RBF-HFD method, the dominant $\mathcal{O}(p^3)$ cost for each ξ_k can also depend on m and n as we use a Greedy algorithm to select the index sets J and \tilde{J} that give weights with desirable properties, as discussed in section 4.4. In practice, $p \ll N$ and would typically be fixed as N increases, so that the total cost scales like $\mathcal{O}(N)$. Furthermore, the weights for one node can be computed independently from the others and is thus an embarrassingly parallel computation. In contrast, the method from [26], requires $\mathcal{O}(N^3)$ operations and results in a dense differentiation matrix. However, the accuracy of this global method is better than the local RBF-HFD approach.




FIG. 1. The maximum absolute value of the residual for differentiating a spherical harmonic, with the maximum taken over the first $n + m$ spherical harmonics. In this figure, the weights are computed using variable precision arithmetic and $\varepsilon = 10^{-10}$. The leveling off of the residual at $\mathcal{O}(10^{-20})$ can be attributed to the choice of ε .

4.3. Choosing the candidate stencil nodes. Increasing the size p of the candidate stencil nodes in the iterated differentiation improves accuracy of the approximations to B^T and C described in the previous section, but also increases the computational cost and worsens the conditioning of the linear system in (19). Theoretically, the smallest possible candidate stencil would simply include every stencil node used in the RBF-HFD formula (11). However, this choice will not lead to accurate weights. Numerical experiments indicate that the candidate stencil should contain at least $n + m$ nodes in order to obtain stable and accurate weights.

In the flat basis limit, as the shape parameter goes to zero, RBF interpolants reproduce certain polynomial interpolants in many cases [11]. Wright and Fornberg [43] provided evidence that the weights obtained from Hermite RBF interpolants in this limit are identical to classical compact weights that are exact for polynomials. For the sphere, it is natural to suppose that the Hermite weights would become exact for the spherical harmonics. This is indeed the case, if the neighborhood size is sufficiently large, which is demonstrated in Figure 1. Let the residual for an example stencil on the sphere be denoted

$$(24) \quad r = \sum_{i=1}^n w_i Y(\mathbf{x}_i) + \sum_{j=1}^m \tilde{w}_j (\Delta_{\mathbb{M}} Y)|_{\mathbf{x}=\tilde{\mathbf{x}}_j} - (\Delta_{\mathbb{M}} Y)|_{\mathbf{x}=\mathbf{x}_1},$$

where Y is a spherical harmonic. Shown in the plot in Figure 1 is the maximum absolute value of the residual, where the maximum is taken over the first $n + m$ spherical harmonics. For $p \ll n + m$, the weights incur errors of very large magnitude, but the error decreases rapidly as p increases. From this plot, we posit that p must consist of at least as many nodes as the number of spherical harmonics of one degree higher than the degree we wish the weights to be exact for. For instance, if we have $n + m = 16$ and we wish the weights to be exact for all spherical harmonics up to third degree (of which there are 16), then p must at least be 25.

Whether the arguments above hold for other surfaces is uncertain, as this depends on the polynomial space spanned by the RBF basis in the limit as ε goes to zero. Additionally, it may not be of particular interest to explore the flat basis limit in practice, as such exploration requires multiple precision arithmetic or a stable method,

such as RBF-GA [30, 20] or RBF-QR [21, 18], for computing the weights. The choice of p should rather be determined by accuracy and stability concerns.

4.4. Greedy algorithm for stencil selection. If the node set is near uniform, experiments have shown that a nearest neighbor approach to stencil selection is usually sound. However, compact stencils can provide additional properties if the stencil nodes are chosen wisely. A simple greedy algorithm, similar to the one in [43], is used for this purpose. For small stencils ($n, m \leq 10$), that provide up to fourth-order convergence, we enforce that all weights in \tilde{L}_Ξ are positive, that \tilde{L}_Ξ is diagonally dominant, and that all off-diagonal elements in L_Ξ are positive. By consistency, any row sum of L_Ξ is zero and thus the diagonal elements are negative. This property of L_Ξ , along with the diagonal dominance of \tilde{L}_Ξ , provides the stability properties outlined in section 6, while imposing positivity of the \tilde{L}_Ξ weights ensures that the compact weights mimic their lattice-based counterparts. For larger stencils, such weights cannot be found, and we must give up the last property.

The greedy algorithm proceeds in the following way:

1. For each node ξ_k , determine the $p - 1$ nearest neighbors to form S_k and compute all matrices necessary to form the approximation to the entries of A_H and the right-hand side in the system (12) as described in section 4.1.
2. Compute all combinations of choosing $n - 1$ nodes from S_k and sort them by average distance to ξ_k . Let $\{J^{(i)}\}_{i=1}^{i_{\max}}$ denote the set of index sets obtained.
3. Repeat step 2 with $n - 1$ replaced by m and denote this set $\{\tilde{J}^{(j)}\}_{j=1}^{j_{\max}}$.
4. Let $i = j = 1$. Until stencils with suitable weights have been found, repeat the following two steps:
5. Compute w and \tilde{w} from the approximation (12) using the stencils $J^{(i)}$ and $\tilde{J}^{(j)}$.
6. If the weights satisfy the conditions, or $i = i_{\max}$ and $j = j_{\max}$, go to step 1 and continue with $k = k + 1$. Else if $j = j_{\max}$, or if $j = 1$ and w is not diagonally dominant, increase i and let $j = 1$. Else increase j .

The rationale behind the second condition in step 6 is that, if w is not diagonally dominant, numerical experiments have shown that replacing the implicit stencil is unlikely to work. For near-uniform nodes and suitable values for the stencil and neighborhood sizes, the conditions are met for $i = j = 1$ for a majority of stencils, which corresponds to the nearest neighbor approach. The algorithm rarely requires more than ten iterations in steps 5 and 6.

5. Using RBF-HFD weights with the method-of-lines. In sections 7.1 and 7.2, we use the RBF-HFD discrete approximation to the surface Laplacian in the method of lines (MOL) to simulate diffusion and reaction-diffusion equations on surfaces. We briefly review this technique for the former equation, as its generalization to the latter follows naturally.

The diffusion of a scalar quantity u on a surface with a (nonlinear) forcing term is given as

$$(25) \quad \frac{\partial u}{\partial t} = \delta \Delta_M u + f(t, u),$$

where $\delta > 0$ is the diffusion coefficient, $f(t, u)$ is the forcing term, and an initial value of u at time $t = 0$ is given.

Our RBF-HFD method for (25) takes the form

$$(26) \quad \frac{d}{dt} u_\Xi = \delta \tilde{L}_\Xi^{-1} L_\Xi u_\Xi + f(t, u_\Xi),$$

where $\tilde{L}_\Xi^{-1}L_\Xi$ is an RBF-HFD discretization of Δ_M over the nodes in Ξ , as described above. This is a system of N coupled ODEs and, provided it is stable (see section 6), can be advanced in time with a suitably chosen time-integration method. In contrast to an explicit RBF-FD discretization, where \tilde{L}_Ξ is the identity matrix, both explicit and implicit time discretizations will require solving a sparse linear system. The diffusion term is typically treated implicitly in order to allow larger time steps, and we have chosen to use a semi-implicit BDF3 method [2], given by

$$(27) \quad \left(\frac{11}{6}I - \delta\Delta t\tilde{L}_\Xi^{-1}L_\Xi\right)u_X^{n+1} = 3u_X^n - \frac{3}{2}u_X^{n-1} + \frac{1}{3}u_X^{n-2} + \Delta t(3f(t^n, u_X^n) - 3f(t^{n-1}, u_X^{n-1}) + f(t^{n-2}, u_X^{n-2})),$$

where the superscript denotes time level. If δ and Δt are constant in time, we may multiply by L_Ξ to obtain the system

$$(28) \quad \underbrace{\left(\frac{11}{6}\tilde{L}_\Xi - \delta\Delta tL_\Xi\right)}_{R_\Xi}u_X^{n+1} = \tilde{L}_\Xi\{\text{r.h.s. of (27)}\}.$$

The matrix R_Ξ is sparse and well-conditioned, and may be factorized if a sufficient amount of memory is available. Another option is to use a Krylov solver with a suitable preconditioner, for instance an ILU decomposition. The latter might be preferable if the time step is adaptive, since the zero-fill ILU preconditioner is cheap to compute, at least compared to a full LU factorization. The performance of this approach is discussed in section 7.

6. Stability and Gershgorin sets. One important reason to favor high-order compact stencils over their explicit counterparts is eigenvalue stability. In the classical setting, where stencils are symmetric, the structure of the obtained generalized eigenvalue problem ensures stability when using any A-stable time stepping scheme. Consider (25) with $f \equiv 0$ and the corresponding compact semidiscretization

$$(29) \quad \frac{d}{dt}u_\Xi = \tilde{L}_\Xi^{-1}L_\Xi u_\Xi.$$

We wish to prove that any eigenvalue of $\tilde{L}_\Xi^{-1}L_\Xi$ lies in the left half-plane, which is a necessary condition for stability. In the following, we will consider the equivalent problem of showing that all eigenvalues of the generalized eigenvalue system

$$(30) \quad A\mathbf{x} = \lambda B\mathbf{x},$$

where $A = -L_\Xi$ and $B = \tilde{L}_\Xi$ lie in the right half-plane.

A common way to prove stability is to use the Gershgorin circle theorem. For instance, if A has zero row sum and positive diagonal and nonpositive off-diagonal elements, then any eigenvalue of A must have a nonnegative real part. We will assume that A has these properties, and that B is a strictly diagonally dominant matrix with positive diagonal elements. If A and B were Hermitian, these properties would be sufficient for the eigenvalues of the generalized eigenvalue problem to have nonnegative real parts. The situation is somewhat more complicated in the nonsymmetric case.

Stewart[41] extended the Gershgorin circle theorem to generalized eigenvalues, and proved that any eigenvalue must lie in $\bigcup_i \Gamma_i$, where Γ_i is given by $z \in \mathbb{C}$ such that

$$(31) \quad |zb_{ii} - a_{ii}| \leq \sum_{j \neq i} |zb_{ij} - a_{ij}|,$$

where a_{ij} and b_{ij} denote the elements of A and B , respectively. In contrast to the regular Gershgorin theorem, it is quite difficult to determine the values of z that fulfill this inequality. By cleverly applying the triangle inequality, Kostić and co-workers [28] provided an approximate Gershgorin set that can be easily computed. Let $r_i(A)$ denote the absolute sum of the i th row of A with the diagonal element zeroed out. The i th approximate Gershgorin set $\hat{\Gamma}_i$ is given by $z \in \mathbb{C}$ such that

$$|zb_{ii} - a_{ii}| \leq |z|r_i(B) + r_i(A).$$

By dividing by b_{ii} , which is positive by assumption, we obtain

$$(32) \quad \left| z - \frac{a_{ii}}{b_{ii}} \right| \leq |z| \frac{r_i(B)}{b_{ii}} + \frac{r_i(A)}{b_{ii}}.$$

We will let $\alpha = \frac{a_{ii}}{b_{ii}}$ and $\beta = \frac{r_i(B)}{b_{ii}}$, and note that we have $r_i(A) = a_{ii}$ from the matrix properties we assumed. Note also that $\beta < 1$ since B is strictly diagonally dominant.

Approximate Gershgorin sets for $\alpha = 1$ and various β are shown in Figure 2. In particular, note that none of the sets contain any part of the negative real axis. For small values of β , the only part of the negative real half-plane that is included in the approximate Gershgorin set is a narrow segment along the imaginary axis. In the special case where B has nonnegative elements, it is typically possible to find stencils that provide $\beta < 0.5$, which makes the unstable part of the Gershgorin set practically insignificant.

In practice, the exact Gershgorin sets turn out to be much smaller than the approximate ones for the discretizations considered here. Examples are shown in Figure 3, where fourth-order and sixth-order approximations of the Laplace–Beltrami operator on the sphere are considered. Note that the regions shown are not the (approximate) Gershgorin set, but rather the i th (approximate) Gershgorin set, where i is chosen as the row that gives the largest extent in the left half-plane. For the fourth-order method, the Gershgorin set shown is essentially completely contained in the right half-plane.




FIG. 2. Approximate Gershgorin sets for $\alpha = 1$ and various values of β . The circle corresponds to $\beta = 0$, and additional curves are given by $\beta = 0.2, 0.4, 0.6, 0.8$, and 0.95 , starting from the innermost to the outermost curve.

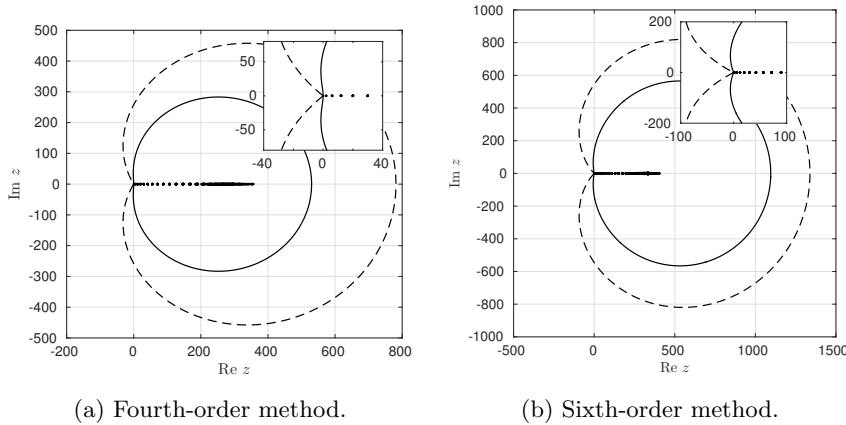


FIG. 3. Generalized eigenvalues of (A, B) and corresponding Gershgorin sets. The i th approximate Gershgorin set is outlined by a dashed line, and the respective exact one is given by a solid line, where the index i is chosen to give the largest extent in the left half-plane of each individual set. The inset shows a magnification about the origin.

TABLE 1

The manifolds in the experiments are given by $\{x, y, z\}$ satisfying $F(x, y, z) = 0$. For the Red Blood Cell, the parameters are $c_0 = \frac{0.81}{a}$, $c_2 = \frac{7.83}{a}$, $c_4 = \frac{-4.39}{a}$, and $r_0 = \frac{3.91}{a}$ with $a = 3.39$. The parameters for Dupin’s Cyclide are $c_1 = 2$, $c_2 = 1.9$, $c_3 = \sqrt{0.39}$, and $c_4 = 1$.

Surface	$F(x, y, z)$
Sphere	$x^2 + y^2 + z^2 - 1$
Torus	$(1 - \sqrt{x^2 + y^2})^2 + z^2 - \frac{1}{9}$
Red Blood Cell	$\left(1 - \frac{x^2 + y^2}{r_0^2}\right) \left(c_0 + c_2 \left(\frac{x^2 + y^2}{r_0^2}\right) + c_4 \left(\frac{x^2 + y^2}{r_0^2}\right)^2\right)^2 - 4z^2$
Dupin’s Cyclide	$(x^2 + y^2 + z^2 - c_4^2 + c_2^2)^2 - 4(c_1x + c_3c_4)^2 - 4c_2^2y^2$
“Tooth”	$x^8 + y^8 + z^8 - (x^2 + y^2 + z^2)$

7. Numerical results. The numerical experiments in this section were performed in MATLAB, using node sets generated by DistMesh [35, 34], unless otherwise noted. All timing experiments were run on a desktop computer equipped with an Intel® Core™2 Quad Q9550 processor at 2.83 GHz and 4 GB of RAM and we did not make explicit use of parallelization in MATLAB. For a mathematical description of the manifolds, see Table 1. Example node sets are presented in Figure 4.

7.1. Parameter studies. To verify the convergence rate and facilitate comparisons between explicit and implicit finite difference methods, we use the forced heat equation with a known analytic solution, given by Eq. (25) with $\delta = 1$. We restrict our attention to the surface of the sphere and the torus, in order to be able to manufacture exact solutions.

As in [38], we take the exact solution for the sphere to be

$$(33) \quad u(t, \mathbf{x}) = e^{-5t} \sum_{k=1}^{23} e^{-10 \cos^{-1}(\mathbf{y}_k \cdot \mathbf{x})},$$




FIG. 4. Example node sets for the red blood cell model, Dupin's cyclide, and the "tooth" model.

where $\mathbf{x} \in \mathbb{S}^2$ and \mathbf{y}_k , $k = 1, \dots, 23$, are randomly placed points on \mathbb{S}^2 . For the torus, we also use the exact solution from [38]:

$$(34) \quad u(t, \lambda, \varphi) = e^{-5t} \sum_{k=1}^{30} e^{-20(1-\cos(\lambda-\lambda_k)) - \frac{9}{4}(1-\cos(\varphi-\varphi_k))}.$$

Here the solution is stated in the parametric variables $(\lambda, \varphi) \in [-\pi, \pi]^2$ for the torus, and (λ_k, φ_k) , $k = 1, \dots, 30$, are taken as randomly chosen values in $[-\pi, \pi]^2$. The exact parameterization of the torus we use is

$$x = \left(1 + \frac{1}{3} \cos(\varphi)\right) \cos(\lambda), \quad y = \left(1 + \frac{1}{3} \cos(\varphi)\right) \sin(\lambda), \quad z = \frac{1}{3} \sin(\varphi).$$

The forcing terms for the diffusion equations on the two surfaces are computed from these exact solutions. Unless otherwise noted, the diffusion equations for both surfaces are simulated for $0 \leq t \leq 0.2$ and the time step is chosen such that spatial errors dominate. All time integrations are done using BDF3 and the forcing function is evaluated implicitly. We restrict the presentation to the relative ℓ_∞ -norm of the error as the observed convergence rates were the same for the ℓ_1 -, ℓ_2 -, and ℓ_∞ -norms. Finally, we let h denote the "spacing" of the nodes in Ξ , and compute this as the average distance to the nearest neighbor.

Shape parameter. Two strategies for the scaling of the shape parameter are commonly used: inversely proportional to the node distance h , or fixed. The former choice keeps the condition number of the regular interpolation matrix, here denoted $\kappa(A_R)$, constant, but introduces a stationary interpolation error that does not decrease to zero in the limit as h goes to zero. A fixed ε gives convergence for all h , but the linear systems for computing the weights become ill-conditioned for small h , and convergence is lost due to round-off errors. There are workarounds for both of these problems. In the stationary interpolation case, it is possible to recover high-order convergence by adding suitable polynomial terms to the interpolant [14, 3]. On a surface, however, the polynomials may themselves introduce ill-conditioning, especially if it is an algebraic surface as polynomial unisolvency becomes an issue.

If ε is kept fixed, there are currently two options to circumvent the problem of ill-conditioning: stable algorithms or variable-precision arithmetic. The algorithms of the former category, such as RBF-GA, RBF-QR, and RBF-CP[20, 18, 22], are unfortunately not easily adaptable to Hermite interpolation in the form introduced here. Instead, we adopt quad-precision arithmetic, for instance using the Advanpix toolbox


[1], which allows accurate determination of the weights for values of h corresponding to millions of nodes on the surfaces considered. Note that quad precision is only needed for the computation of the weights, after which the results can be truncated back to double precision and the simulations run without issues. Also note that computing the weights is an embarrassingly parallel task, and one that is only performed once for a given simulation. Optimization of ε is beyond the scope of this study, and the value is chosen such that weights can be computed in double precision for small to moderate node sets ($N \lesssim 50,000$), after which we switch to quad precision.

Results for the forced heat equation of the sphere illustrating the effects of the two shape parameter strategies are presented in Figures 5(a) and 5(b), in which the errors as a function of h are shown for a fourth-order and a sixth-order approximation, respectively. In the fourth-order case, both strategies for selecting the shape parameter converge with the same rate for a large range of values of h , but keeping the condition number fixed eventually results in a stationary error as h decreases. For the sixth-order stencil, the resulting order is slightly lower in the fixed condition number setting, and convergence is achieved only in a small range of values of h . Note that the mean condition number of the Hermite interpolation matrix, A_H , was kept fixed in these experiments. Similar results for the torus were observed and have thus been omitted.


Stencil sizes. The choice $m = n - 1$ and letting the index sets I and J coincide (bar the evaluation node) provides ideal sparsity of the matrix R_{Ξ} in Eq. (28). However, as the approximation order is increased, it gradually becomes more difficult to find weights w that satisfy the diagonal dominance criterion. On near-uniform nodes, it is typically possible to find stable weights for n up to 12. The stencil selection algorithm will also influence the sparsity of the matrices, and certain combinations of n and m are more likely to produce stable weights (e.g., by symmetry). Another consideration is the formal approximation order of the stencil, which we are not able to derive, and so instead determine it experimentally. The order appears to be primarily determined by the number of degrees of freedom of the Hermite system, i.e., $n + m$.

Figures 6(a) and 6(b) show the error as a function of h for the forced heat equation on the sphere and on the torus, respectively, with different values of m and n . In these figures, the shape parameter is kept fixed at $\varepsilon = 3$ for the smaller stencil sizes and $\varepsilon = 5$ for the larger ones. With these choices, quad-precision arithmetic is only required for computing the weights for the two largest node sets, which range in size from $N \approx 1000$ to $N \approx 200,000$. We summarize the observations from this experiment regarding observed order of convergence and stencil sizes in Table 2.

Performance comparison. In addition to the stability properties discussed in section 6, compact stencils provide better sparsity for the same approximation order. This should lead to a smaller memory footprint and fewer floating-point operations for solving the system in Eq. (28). In this paper, we use BiCGSTAB with zero-fill ILU preconditioner for solving the linear system. Table 3 shows the number of nonzero elements in R_{Ξ} and the CPU time for a simulation for some combinations of N , n , and m . Also presented in this table is the average number of Krylov iterations per time step. It is interesting to note from this table that, in terms of CPU time, the smaller implicit stencil barely outperforms the explicit stencil of the same order. This can be attributed to the larger number of iterations needed for the Krylov solver to converge. Increasing the stencil size to $n = 12$ and $m = 15$ reduces the number of iterations needed, plausibly due to the larger number of nonzeros in the incomplete LU factorization. This results in a CPU time that is comparable to that of the smaller implicit stencil. In this comparison, the fourth-order explicit stencil ($n = 17$, $m = 0$)




(a) $n = 10, m = 6, p = 19$.




(b) $n = 12, m = 15, p = 32$.

FIG. 5. The error as a function of h for the forced heat equation on the sphere for different strategies of shape parameter selection. Solid lines correspond to fixed shape parameter, and dashed lines correspond to fixed condition number. The dash-dot line show slopes for convergence rate $\mathcal{O}(h^p)$ with $p = 4, 6$. In (a) the values of ε are $\{6, 5, 4, 3, 2.5\}$ from top to bottom, and in (b) the values are $\{8, 7, 6, 5, 4.5\}$, again from top to bottom. The values of $\kappa(A_H)$ are $\{10^{10}, 10^{11}, 10^{12}, 10^{13}, 10^{14}\}$ from top to bottom (at small h) in both (a) and (b).



(a) Sphere.



(b) Torus.

FIG. 6. The error as a function of h for the forced heat equation with different stencil sizes. The dashed lines show slopes for convergence rate $\mathcal{O}(h^p)$, fitted from the data points.

TABLE 2

Recommended stencil and neighborhood sizes for different approximation orders. The order was determined from numerical experiments on the sphere and the torus shown in Figure 6.

n	m	p	Observed order
10	6	19	4
12	15	32	6
17	0	17	4
31	0	31	6

TABLE 3

The table shows the number of nonzeros of the matrix R_{Ξ} , the average iteration count for the Krylov solver, and the CPU time for 200 time steps with $\Delta t = 10^{-3}$ for some choices of n , m , and h . The surface in this experiment is the sphere, and the tolerance for the Krylov solver is 10^{-12} .

n	m	h	N	# of nonzeros	Avg. iter.	CPU time (s)
10	6	0.1	1806	18064	2	0.37
		0.05	7446	74462	3	1.5
		0.025	30054	300548	5.5	11
12	15	0.1	1806	28925	2	0.41
		0.05	7446	119136	2.5	1.7
		0.025	30054	480864	4	10
17	0	0.1	1806	30702	2	0.39
		0.05	7446	126582	2.5	1.6
		0.025	30054	510918	4	10
31	0	0.1	1806	55986	1.5	0.40
		0.05	7446	230826	2	2.4
		0.025	30054	931674	4	18




FIG. 7. The run-time for a simulation of the forced heat equation on the sphere as a function of h , using $\Delta t = 10^{-2} \cdot h$. The simulation was run from $t = 0$ to $t = 0.2$.

is on par with the sixth-order implicit stencil both in terms of memory requirements and computational cost. Note that the time step chosen, $\Delta t = 10^{-3}$, is rather small (which is needed for spatial errors to dominate). Increasing the time step to $\Delta t = 0.1$ increases the number of Krylov iterations per time step roughly by a factor of 6.

In most applications, Δt would be chosen proportional to h in order to reduce both spatial and temporal errors as the node set is refined. Figure 7 shows the CPU

time as a function of h with $\Delta t = 10^{-2} \cdot h$. For the near-uniform node sets used throughout this paper, $N \propto 1/\sqrt{h}$, and thus the CPU time scales as $\mathcal{O}(N^{3/2})$.

7.2. Applications. We now present applications of the new compact scheme to solving reaction-diffusion equations on different surfaces. As in [38], we present results of simulations both on surfaces defined implicitly by algebraic expressions (see Table 1) and on more general point cloud surfaces. On the former, we simulate the same two-species Turing system used in [38], given by

$$(35) \quad \frac{\partial u}{\partial t} = \alpha u(1 - \tau_1 v^2) + v(1 - \tau_2 u) + \delta_u \Delta_{\mathbb{M}} u,$$

$$(36) \quad \frac{\partial v}{\partial t} = \beta v \left(1 + \frac{\alpha \tau_1}{\beta} uv \right) + u(\gamma + \tau_2 v) + \delta_v \Delta_{\mathbb{M}} v.$$

A visualization of the solutions of this equation on the various surfaces with the parameters selected from Table 4 are shown in Figure 8.

TABLE 4

The table shows the values of the parameters of (35) and (36) used in the numerical experiments shown in Figure 8. We set $\delta_u = 0.516\delta_v$ for the red blood cell and tooth models, and use $\delta_u = 5.16\delta_v$ for Dupin’s cyclide.

Pattern	δ_v	α	β	γ	τ_1	τ_2	Final time
Spots	4.5×10^{-3}	0.899	-0.91	-0.899	0.02	0.2	200
Stripes	2.1×10^{-3}	0.899	-0.91	-0.899	3.5	0	4000

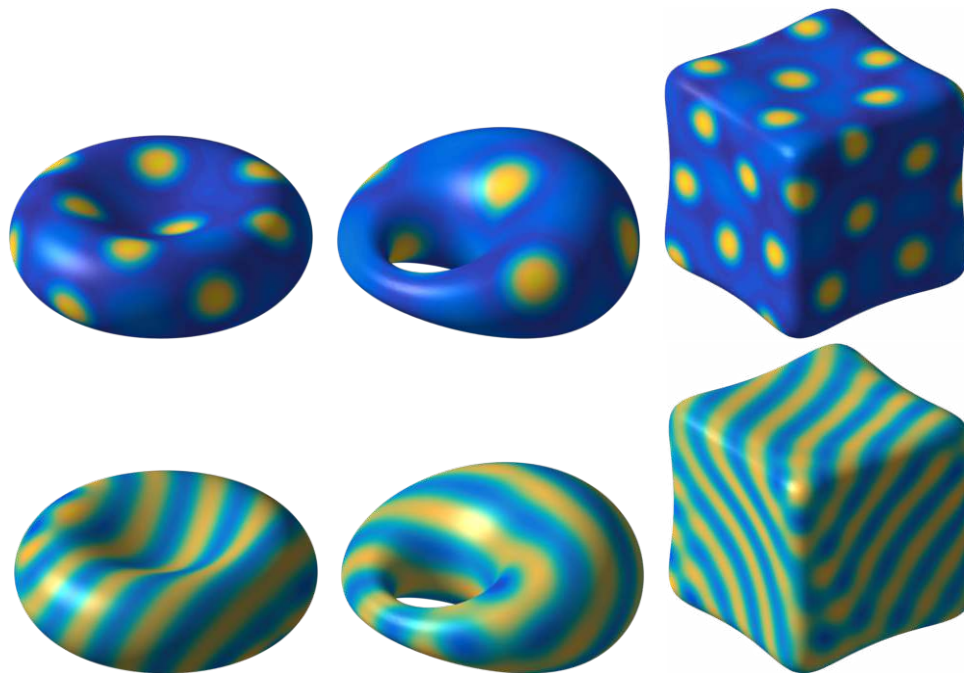


FIG. 8. Quasi-steady Turing spots and stripe patterns resulting from solving (35) and (36) on the red blood cell model, Dupin’s cyclide, and the “tooth” model using the fourth-order compact scheme. In all plots, light shading indicates high concentrations and dark shading indicates low concentrations (the electronic version is full-color).




FIG. 9. Fitzhugh–Nagumo spiral wave patterns resulting from solving (37) and (38) on the bumpy sphere and bunny models using the fourth-order compact scheme. In all plots, light shading indicates high concentrations and dark shading indicates low concentrations (the electronic version is full-color).

On the more general point cloud surfaces, we simulate the Fitzhugh–Nagumo-type model used in [26]:

$$(37) \quad \frac{\partial u}{\partial t} = \delta_u \Delta_{\mathbb{M}} u + \frac{1}{\alpha} u(1-u) \left(u - \frac{v+b}{a} \right),$$

$$(38) \quad \frac{\partial v}{\partial t} = \delta_v \Delta_{\mathbb{M}} v + u - v.$$

For both the bumpy sphere² and bunny³ models, we set $a = 0.75$, $b = \alpha = 0.02$, $\delta_v = 0$, and $\delta_u = 1.5(\frac{2\pi}{50})^2$. With these parameters, the model generates dynamic spiral wave solutions. Snapshots of these solutions computed with the RBF-HFD method for the two surfaces are shown in Figure 9. We note that the normal vectors on these point cloud models can be generated by any appropriate method. In this work, the node sets and normal vectors were created using MeshLab [8], utilizing the Poisson surface reconstruction algorithm with some additional smoothing and the Poisson disk sampling method.

7.3. Curvature, node spacing, and the stencil selection algorithm. In a few of our experiments, the greedy algorithm failed to find suitable stencils for some

²Available from the Aim@Shape Shape Repository (<http://visionair.ge.imati.cnr.it/>).

³Available from the Stanford Computer Graphics Laboratory (<http://graphics.stanford.edu/data/3Dscanrep/>).




FIG. 10. The number of rejected stencils as a function of h and the mean curvature H .

node points. A common characteristic of these nodes was that they were located in areas of large curvature, e.g., around the ears of the bunny. To illustrate this curvature effect more clearly on the failure of the stencil selection algorithm, we carried out an experiment generating the surface Laplacian on prolate spheroids of varying curvatures and determining how the number of rejected stencils changed with curvature. The results are plotted in Figure 10 in terms of the number of rejected stencils, i.e., stencils where the stability constraints could not be met, as a function of both the node distance h and the maximum mean curvature H . Note that i_{\max} and j_{\max} in the greedy algorithm were set to 5 in this experiment so that only a small number of stencils were attempted for each node, to emphasize the effect of the curvature. This experiment indicates that decreasing the node spacing alleviates the issue. In our experiments on the bunny, we found that simply using nearest neighbor stencils for cases where no suitable stencil could be found still led to temporal stability, and we therefore recommend this approach. Two other potential remedies are increasing ε for the nodes where no suitable stencils can be found, and refining the node set locally in these areas. The former has previously been noted to improve stability (see, e.g., [16]), but larger values of ε also tend to reduce the accuracy. A full exploration of these latter two approaches is beyond the scope of this work.

8. Discussion. The compact RBF-HFD scheme improves on previous RBF-FD schemes for diffusion on surfaces both in terms of efficiency and stability. The proposed greedy stencil selection algorithm ensures eigenvalue stability on surfaces without large (or rapid) changes in curvature. In numerical experiments of the forced diffusion equation on the sphere and the torus, the new scheme provided accuracy similar to the previous noncompact RBF-FD method, but with a smaller memory footprint and higher accuracy. The linear systems generated from the semi-implicit BDF3 time discretization were also shown to be efficiently solvable using BiCGStab with a standard zero-fill ILU preconditioner. In addition to illustrating the good stability, accuracy, and efficiency properties of the scheme, we showed how it can be easily adaptable to reaction-diffusion equations on both implicitly defined surfaces and surfaces defined by a point cloud. The method can be naturally generalized to a smooth orientable surface that is discretized with a set of roughly uniform nodes and with approximations to the normal to the surface at each of the nodes.

While stencil sizes generating fourth- and sixth-order convergence in numerical experiments on the sphere and the torus are provided, no investigation of the theoretical convergence rates have been given. This is clearly an avenue of future investigation. Another future topic of research is the influence of the curvature and the shape parameter on the computed RBF-HFD weights. Experiments suggest that large local curvature makes it impossible to find weights that satisfy the conditions that ensure eigenvalue stability, although no issues with temporal integration were encountered when this stability was not insured. Refining the node sets in areas of high curvature or increasing the shape parameters used in the stencils in these areas may provide a simple way to alleviate the problem. An extensive investigation of the relationship between curvature, nodal distance, and stability will be the topic of a future study.

Finally, we note that an extension to convection-diffusion problems would allow the method to be used in various applications, e.g., chemical transport on thin membranes and shells, biomechanical modeling of cells. This is currently being pursued by the first author.

REFERENCES

- [1] *Advantix Multiprecision Computing Toolbox for MATLAB*, <http://www.advantix.com/>, accessed 2016-06-09.
- [2] U. M. ASCHER, S. J. RUUTH, AND B. T. R. WETTON, *Implicit-explicit methods for time-dependent PDEs*, SIAM J. Numer. Anal. 32 (1997), pp. 797–823.
- [3] V. BAYONA, N. FLYER, B. FORNBERG, AND G. A. BARNETT, *On the role of polynomials in RBF-FD approximations: II. Numerical solution of elliptic PDEs*, J. Comput. Phys., 332 (2017), pp. 257–273.
- [4] V. BAYONA, M. MOSCOSO, M. CARRETERO, AND M. KINDELAN, *RBF-FD formulas and convergence properties*, J. Comput. Phys., 229 (2010), pp. 8281–8295.
- [5] R. L. BISHOP AND S. I. GOLDBERG, *Tensor Analysis on Manifolds*, Macmillan, New York, 1968.
- [6] T. CECIL, J. QIAN, AND S. OSHER, *Numerical methods for high dimensional Hamilton–Jacobi equations using radial basis functions*, J. Comput. Phys., 196 (2004), pp. 327–347.
- [7] G. CHANDHINI AND Y. SANYASIRAJU, *Local RBF-FD solutions for steady convection–diffusion problems*, Int. J. Numer. Meth., 72 (2007), pp. 352–378.
- [8] P. CIGNONI, M. CORSINI, AND G. RANZUGLIA, *MESHLAB: An open-source 3D mesh processing system*, ERCIM News, (2008), pp. 45–46.
- [9] L. COLLATZ, *The Numerical Treatment of Differential Equations*, 3rd ed., Springer, Berlin, 1966.
- [10] O. DAVYDOV AND D. T. OANH, *Adaptive meshless centres and RBF stencils for Poisson equation*, J. Comput. Phys., 230 (2011), pp. 287–304.
- [11] T. A. DRISCOLL AND B. FORNBERG, *Interpolation in the limit of increasingly flat radial basis functions*, Comput. Math. Appl., 43 (2002), pp. 413–422.
- [12] G. E. FASSHAUER, *Meshfree Approximation Methods with MATLAB*, Interdisciplinary Mathematical Sciences 6, World Scientific, Singapore, 2007.
- [13] G. E. FASSHAUER AND M. J. MCCOURT, *Stable evaluation of Gaussian radial basis function interpolants*, SIAM J. Sci. Comput., 34 (2012), pp. A737–A762.
- [14] N. FLYER, B. FORNBERG, V. BAYONA, AND G. A. BARNETT, *On the role of polynomials in RBF-FD approximations I. Interpolation and accuracy*, J. Comput. Phys., 321 (2016), pp. 21–38.
- [15] N. FLYER, E. LEHTO, S. BLAISE, G. B. WRIGHT, AND A. ST-CYR, *A guide to RBF-generated finite differences for nonlinear transport: Shallow water simulations on a sphere*, J. Comput. Phys., 231 (2012), pp. 4078–4095.
- [16] N. FLYER AND G. B. WRIGHT, *A radial basis function method for the shallow water equations on a sphere*, Proc. R. Soc. A, 465 (2009), pp. 1949–1976.
- [17] B. FORNBERG AND N. FLYER, *A Primer on Radial Basis Functions with Applications to the Geosciences*, SIAM, Philadelphia, 2014.
- [18] B. FORNBERG, E. LARSSON, AND N. FLYER, *Stable computations with Gaussian radial basis functions*, SIAM J. Sci. Comput., 33 (2011), pp. 869–892.
- [19] B. FORNBERG AND E. LEHTO, *Stabilization of RBF-generated finite difference methods for convective PDEs*, J. Comput. Phys., 230 (2011), pp. 2270–2285.

- [20] B. FORNBERG, E. LEHTO, AND C. POWELL, *Stable calculation of Gaussian-based RBF-FD stencils*, *Comput. Math. Appl.*, 65 (2013), pp. 627–637.
- [21] B. FORNBERG AND C. PIRET, *A stable algorithm for flat radial basis functions on a sphere*, *SIAM J. Sci. Comput.*, 30 (2007), pp. 60–80.
- [22] B. FORNBERG AND G. WRIGHT, *Stable computation of multiquadric interpolants for all values of the shape parameter*, *Comput. Math. Appl.*, 48 (2004), pp. 853–867.
- [23] B. FORNBERG AND J. ZUEV, *The Runge phenomenon and spatially variable shape parameters in RBF interpolation*, *Comput. Math. Appl.*, 54 (2007), pp. 379–398.
- [24] E. FUSELIER AND G. WRIGHT, *Scattered data interpolation on embedded submanifolds with restricted positive definite kernels: Sobolev error estimates*, *SIAM J. Numer. Anal.*, 50 (2012), pp. 1753–1776.
- [25] E. FUSELIER AND G. B. WRIGHT, *Order-preserving derivative approximation with periodic radial basis functions*, *SIAM J. Numer. Anal.*, 41 (2015), pp. 23–53.
- [26] E. J. FUSELIER AND G. B. WRIGHT, *A high-order kernel method for diffusion and reaction-diffusion equations on surfaces*, *J. Sci. Comput.*, (2013), pp. 1–31.
- [27] Q. T. L. GIA, *Approximation of parabolic pdes on spheres using spherical basis functions*, *Adv. Comput. Math.*, 22 (2005), pp. 377–397.
- [28] V. KOSTIĆ, R. S. VARGA, AND L. CVETKOVIĆ, *Localization of generalized eigenvalues by Cartesian ovals*, *Numer. Linear Algebra Appl.*, 19 (2012), pp. 728–741.
- [29] E. LARSSON AND B. FORNBERG, *Theoretical and computational aspects of multivariate interpolation with increasingly flat radial basis functions*, *Comput. Math. Appl.*, 49 (2005), pp. 103–130.
- [30] E. LARSSON, E. LEHTO, A. HERYUDONO, AND B. FORNBERG, *Stable computation of differentiation matrices and scattered node stencils based on Gaussian radial basis functions*, *SIAM J. Sci. Comput.*, 35 (2013), pp. A2096–A2119.
- [31] S. K. LELE, *Compact finite difference schemes with spectral-like resolution*, *J. Comput. Phys.*, 103 (1992), pp. 16–42.
- [32] C. MACDONALD AND S. RUUTH, *The implicit closest point method for the numerical solution of partial differential equations on surfaces*, *SIAM J. Sci. Comput.*, 31 (2010), pp. 4330–4350.
- [33] F. NARCOWICH AND J. WARD, *Generalized hermite interpolation via matrix-valued conditionally positive definite functions*, *Math. Comput.*, 63 (1994), pp. 661–688.
- [34] P.-O. PERSSON, *DistMesh – a simple mesh generator in MATLAB*. <http://persson.berkeley.edu/distmesh/>, accessed 2016-06-09.
- [35] P.-O. PERSSON AND G. STRANG, *A simple mesh generator in MATLAB*, *SIAM Rev.*, 46 (2004), pp. 329–345.
- [36] C. PIRET, *The orthogonal gradients method: A radial basis functions method for solving partial differential equations on arbitrary surfaces*, *J. Comput. Phys.*, 231 (2012), pp. 4662–4675.
- [37] R. SCHABACK, *Multivariate interpolation by polynomials and radial basis functions*, *Constr. Approx.*, 21 (2005), pp. 293–317.
- [38] V. SHANKAR, G. B. WRIGHT, R. M. KIRBY, AND A. L. FOGELSON, *A radial basis function (RBF)-finite difference (FD) method for diffusion and reaction-diffusion equations on surfaces*, *J. Sci. Comput.*, 63 (2014), pp. 745–768.
- [39] C. SHU, H. DING, AND K. S. YEO, *Local radial basis function-based differential quadrature method and its application to solve two-dimensional incompressible Navier–Stokes equations*, *Comput. Methods Appl. Mech. Eng.*, 192 (2003), pp. 941–954.
- [40] D. STEVENS, H. POWER, M. LEES, AND H. MORVAN, *The use of PDE centers in the local RBF Hermitean method for 3D convective-diffusion problems*, *J. Comput. Phys.*, 228 (2009), pp. 4606–4624.
- [41] G. W. STEWART, *Gershgorin theory for the generalized eigenvalue problem $Ax = \lambda Bx$* , *Math. Comput.*, 29 (1975), pp. 600–606.
- [42] G. B. WRIGHT, N. FLYER, AND D. A. YUEN, *A hybrid radial basis function–pseudospectral method for thermal convection in a 3-D spherical shell*, *Geochem. Geophys. Geosyst.*, 11 (2010).
- [43] G. B. WRIGHT AND B. FORNBERG, *Scattered node compact finite difference-type formulas generated from radial basis functions*, *J. Comput. Phys.*, 212 (2006), pp. 99–123.
- [44] W. ZONGMIN, *Hermite–Birkhoff interpolation of scattered data by radial basis functions*, *Approx. Theory Appl.*, 8 (1992), pp. 1–10.