

11-15-2021

An Efficient High-Order Meshless Method for Advection-Diffusion Equations on Time-Varying Irregular Domains

Varun Shankar
University of Utah

Grady B. Wright
Boise State University

Aaron L. Fogelson
University of Utah

Publication Information

Shankar, Varun; Wright, Grady B.; and Fogelson, Aaron L. (2021). "An Efficient High-Order Meshless Method for Advection-Diffusion Equations on Time-Varying Irregular Domains". *Journal of Computational Physics*, 445, 110633. <https://doi.org/10.1016/j.jcp.2021.110633>

This is an author-produced, peer-reviewed version of this article. © 2021, Elsevier. Licensed under the Creative Commons Attribution-NonCommercial-No Derivative Works 4.0 International license. The final, definitive version of this document can be found online at *Journal of Computational Physics*, <https://doi.org/10.1016/j.jcp.2021.110633>

An Efficient High-Order Meshless Method for Advection-Diffusion Equations on Time-Varying Irregular Domains

Varun Shankar^{a,*}, Grady B. Wright^b, Aaron L. Fogelson^c

^a*School of Computing, University of Utah, UT, USA*

^b*Department of Mathematics, Boise State University, ID, USA*

^c*Departments of Mathematics and Biomedical Engineering, University of Utah, UT, USA*

Abstract

We present a high-order radial basis function finite difference (RBF-FD) framework for the solution of advection-diffusion equations on time-varying domains. Our framework is based on a generalization of the recently developed Overlapped RBF-FD method that utilizes a novel automatic procedure for computing RBF-FD weights on stencils in variable-sized regions around stencil centers. This procedure eliminates the overlap parameter δ , thereby enabling tuning-free assembly of RBF-FD differentiation matrices on moving domains. In addition, our framework utilizes a simple and efficient procedure for *updating* differentiation matrices on moving domains tiled by node sets of time-varying cardinality. Finally, advection-diffusion in time-varying domains is handled through a combination of rapid node set modification, a new high-order semi-Lagrangian method that utilizes the new tuning-free overlapped RBF-FD method, and a high-order time-integration method. The resulting framework has no tuning parameters and has $O(N \log N)$ time complexity. We demonstrate high-orders of convergence for advection-diffusion equations on time-varying 2D and 3D domains for both small and large Peclet numbers. We also present timings that verify our complexity estimates. Finally, we utilize our method to solve a coupled 3D problem motivated by models of platelet aggregation and coagulation, once again demonstrating high-order convergence rates on a moving domain.

Keywords: Radial basis function; high-order method; meshfree; advection-diffusion; RBF-FD; semi-Lagrangian.

1. Introduction

Collocation methods based on radial basis functions (RBFs) have been increasingly popular for numerically solving partial differential equations (PDEs), due to their high-order convergence rates and their ability to naturally handle scattered node layouts on arbitrary domains. RBF interpolants can be used to generate both pseudospectral (RBF-PS) and finite-difference (RBF-FD) methods [2, 6, 15, 23, 24, 67]. RBF-based methods are also easily applied to the solution of PDEs on node sets that are not unisolvent for polynomials, such as ones lying on the sphere \mathbb{S}^2 [25–28] and other general surfaces [29, 36, 46, 47, 55].

The focus of this paper is on advection-diffusion problems on domains $\Omega(t)$ with boundary conditions enforced at time-varying internal embedded boundaries and a fixed outer boundary. This can be modeled by the

*Corresponding Author

Email addresses: shankar@cs.utah.edu (Varun Shankar), gradywright@boisestate.edu (Grady B. Wright), fogelson@math.utah.edu (Aaron L. Fogelson)

following equations:

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

$$\alpha(\mathbf{x}, t) \mathbf{n} \cdot \nabla c + \beta(\mathbf{x}, t) c = g(\mathbf{x}, t), \mathbf{x} \in \partial\Omega(t), \quad (2)$$

where $c(\mathbf{x}, t)$ is a scalar quantity being transported in the incompressible velocity field $\mathbf{u}(\mathbf{x}, t)$, ν is the diffusion coefficient, α and β are functions that determine boundary conditions (linear in this article), \mathbf{n} are the unit *outward* normals to the domain, and $g(\mathbf{x}, t)$ is either a prescribed or numerically computed boundary condition. Our interest in the above equations stems from their application in the modeling and simulation of platelet aggregation and coagulation [37–39]. Broadly speaking, numerical methods to solve such systems can be divided into three categories: (a) **Eulerian** methods (b) **Lagrangian** methods and (c) **semi-Lagrangian** (SL) methods.

There is extensive literature on Eulerian finite difference (FD) or finite volume (FV) methods for solving PDEs along with boundary conditions on fixed irregular surfaces embedded in the computational domain. Such methods (most of which are designed for Cartesian grids) are mainly of two types. The first type involves *augmenting* the FD/FV scheme to enforce boundary conditions at the irregular boundary. This could be done via *spreading* and *restriction* as in the immersed boundary (IB) method [43–45], or via adding unknowns to the system to force the PDE to satisfy boundary conditions (at the irregular embedded boundaries) as in the wide class of *forcing methods* [31, 34, 42, 54, 69], the ghost cell method [14, 30], and the more recent immersed boundary smooth extension method [62, 63]. In contrast, the second type involves modifying FD/FV stencils near the boundary, such as in the original direct forcing method [19], the immersed interface method (IIM) [40], the embedded boundary method (EBM) [33], the sharp interface method [65, 70], and the capacity function finite volume method [13]. Finally, to tackle moving boundaries outside the original IB framework, a common approach involves *converting* the moving-boundary problem into a series of fixed-boundary problems each solved by one of the above approaches (*e.g.*, see [41]). Such methods typically require an additional spatial extrapolation step to fill newly-uncovered grid points as the domain boundary moves. For all these types of Eulerian methods, obtaining a stable, high-order discretization in space and time can be challenging both due to the presence of a background Cartesian grid and the need for spatial extrapolation to fill newly-uncovered grid points.

In contrast to Eulerian methods that use a fixed background grid, Lagrangian methods involve populating the moving domain with a set of marker particles that move with the velocity field \mathbf{u} . In this case, the advection term is handled without any difficulty. However, to discretize the diffusion term, one of the following approaches can be used: (a) interpolate quantities to a fixed background grid and discretize the PDE there like in the material point method (MPM) [32]; (b) discretize the diffusion term directly on the distorted Lagrangian grid as in smoothed particle hydrodynamics (SPH) [64]; or (c) some variation on the weighted particle method [17] (sometimes called the particle strength exchange method). In all cases, it may be necessary to introduce some form of Lagrangian particle rearrangement to improve spatial resolution and convergence rates [12]. Alternatively, it is possible to reformulate the PDE using the Feynman-Kac formula so that the diffusion term is also handled in a Lagrangian fashion [10, 11, 21].

Semi-Lagrangian (SL) methods are Eulerian methods that use fictitious Lagrangian particles to determine the numerical domain of dependence. We focus on the class of *backward* SL methods, which have found wide application to problems in fluid dynamics, climate modeling, and numerical weather prediction [58–61, 68]. For a pure advection equation, these methods assume that Lagrangian marker particles have arrived at every time-step on an Eulerian grid (or more generally node set). By tracing these particles backward through the velocity field (and to a previous time level), determining their departure positions, and interpolating the solution to those departure positions from the fixed Eulerian node set, one can determine how much material was advected to a given Eulerian location. Solving an advection-diffusion problem then amounts to using an appropriate splitting scheme. The advantage of this method is that the diffusion operator is always discretized on the Eulerian grid. In addition, when solving problems with moving boundaries, the SL framework obviates the need for any spatial extrapolation.

The numerical method presented in this paper relies on the SL framework for precisely these reasons. While RBF methods have been used within the SL framework before, these either were global RBF methods [35],

methods that relied on Voronoi cells [7], or localized RBF methods designed specifically for the sphere [53]. Our new method is based on a generalization of the overlapped RBF-FD method [49, 50, 52, 57], and therefore allows the use of scattered or quasi-uniform nodes in place of a background Cartesian grid, allowing for arbitrary outer boundaries. The time-varying nature of the domain is handled by enabling or disabling pre-existing background nodes contained by the moving boundaries (and in a *small* neighborhood around them). To facilitate this node set adaptation, we represent the moving boundaries using a high-order accurate parametric model built from Lagrangian markers. In this way, our node sets always conform to the time-varying domain, unlike the Eulerian methods discussed above. This technique allows efficient updates to differentiation matrices and necessitates recalculation of overlapped RBF-FD weights *only* in neighborhoods around the moving boundaries. The resulting meshless method allows for high orders of spatial and temporal convergence, does not require spatial extrapolation, and is of quasi-linear computational complexity.

The remainder of the paper is organized as follows. In Section 2, we present our generalization of the overlapped RBF-FD method that removes tuning parameters. In Section 3, we present and describe our overall numerical method in Algorithm 2, complete with error estimates and parameter choices. We present a simple and efficient preconditioner in Section 4, which we then use to solve the time-varying sparse linear system resulting from our numerical method. We conduct a thorough complexity analysis of Algorithm 2 in Section 5. Then, in Section 6, we present 2D and 3D convergence tests on problems with moving embedded boundaries for a range of Peclet numbers. Finally, in Section 7, we present an application of our method to a 3D coupled problem with time-varying boundary conditions on a moving domain inspired by mathematical models of platelet aggregation and coagulation. We conclude with a summary and comments on future work in Section 8.

2. An automatic overlapped RBF-FD method

We now present a generalization of the overlapped RBF-FD method that eliminates the overlap parameter δ , and instead automatically computes, tests, and retains/discards candidate weights on a given stencil. Our approach for automation is to use two stability indicators to indicate whether a set of computed weights is of sufficient quality. In the discussion that follows, we will primarily focus on the new method, remarking on the older version presented in [49, 50, 52] as needed.

Let $X = \{\mathbf{x}_k\}_{k=1}^N$ be a global set of nodes on a domain $\Omega \subset \mathbb{R}^d$. Define the stencil P_k to be the set of nodes containing node $\mathbf{x}_{\mathcal{I}_1^k}$ and its $n-1$ nearest neighbors $\{\mathbf{x}_{\mathcal{I}_2^k}, \dots, \mathbf{x}_{\mathcal{I}_n^k}\}$; here, $\{\mathcal{I}_1^k, \dots, \mathcal{I}_n^k\}$ are indices that map into the global node set X and $\mathcal{I}_1^k = k$. Without loss of generality, we focus on the stencil P_1 . Let $1 \leq p_1 \leq n$ be the number of points on this stencil for which we wish to compute RBF-FD weights. Further, define R_1 to be the global indices of these p_1 nodes so that

$$R_1 = \{\mathcal{R}_1^1, \mathcal{R}_2^1, \dots, \mathcal{R}_{p_1}^1\}. \quad (3)$$

Next, let \mathbb{B}_1 be the ball containing the nodes whose indices are in R_1 . Thus,

$$\mathbb{B}_1 = \{\mathbf{x}_{\mathcal{R}_1^1}, \dots, \mathbf{x}_{\mathcal{R}_{p_1}^1}\}. \quad (4)$$

We discuss how to obtain \mathbb{B}_1 and R_1 in Section 2.2. First, in Section 2.1, we describe how to compute RBF-FD weights for all the nodes in the ball \mathbb{B}_1 .

2.1. Computing weights

The weights for all the nodes in \mathbb{B}_1 with indices in R_1 are computed using the following augmented local RBF interpolant on P_1 :

$$s_1(\mathbf{x}, \mathbf{y}) = \sum_{j=1}^n w_j^1(\mathbf{y}) \|\mathbf{x} - \mathbf{x}_{\mathcal{I}_j^1}\|^m + \sum_{i=1}^M \lambda_i^1(\mathbf{y}) \psi_i^1(\mathbf{x}), \quad (5)$$

where $\|\mathbf{x} - \mathbf{x}_{\mathcal{I}_j^1}\|^m$ is the polyharmonic spline (PHS) RBF of degree m (m is odd), and $\{\psi_i^1(\mathbf{x})\}$ form a basis for the space of polynomials of total degree degree ℓ in d dimensions so that $M = \binom{\ell+d}{d}$; common

choices for these include monomials [50] or orthogonal polynomials [52]. In this work, we select the $\psi_i^1(\mathbf{x})$ functions to be Legendre polynomials. The n overlapped RBF-FD weights associated with the point \mathbf{y} are $w_j^1(\mathbf{y}), j = 1, \dots, n$. We compute the weights for the linear operator \mathcal{L} uniquely at all nodes in \mathbb{B}_1 with indices in the set R_1 by imposing the following two (sets of) conditions on (5):

$$s_1|_{\mathbf{x} \in P_1, \mathbf{y} \in \mathbb{B}_1} = \mathcal{L}\|\mathbf{x} - \mathbf{x}_{\mathcal{I}_j^1}\|^m|_{\mathbf{x} \in \mathbb{B}_1}, j = 1, \dots, n, \quad (6)$$

$$\sum_{j=1}^n w_j^1(\mathbf{y})\psi_i^1(\mathbf{x})|_{\mathbf{x} \in P_1, \mathbf{y} \in \mathbb{B}_1} = \mathcal{L}\psi_i^1(\mathbf{x})|_{\mathbf{x} \in \mathbb{B}_1}, i = 1, \dots, M. \quad (7)$$

These conditions enforce that the weights are exact for both \mathcal{L} applied to the PHS RBF and to the polynomial basis. In this work, we use the heuristic $n = 2M + 1$ [23, 24, 49, 50]; however, larger stencil sizes can sometimes be beneficial [4]. The constraints (6)–(7) for determining the weights in (5) can be collected into the following block linear system:

$$\begin{bmatrix} A_1 & \Psi_1 \\ \Psi_1^T & 0 \end{bmatrix} \begin{bmatrix} W_1 \\ W_1^\psi \end{bmatrix} = \begin{bmatrix} B_{A_1} \\ B_{\Psi_1} \end{bmatrix}, \quad (8)$$

where

$$(A_1)_{ij} = \|\mathbf{x}_{\mathcal{I}_i^1} - \mathbf{x}_{\mathcal{I}_j^1}\|^m, i, j = 1, \dots, n, \quad (9)$$

$$(\Psi_1)_{ij} = \psi_j^1(\mathbf{x}_{\mathcal{I}_i^1}), i = 1, \dots, n, j = 1, \dots, M, \quad (10)$$

$$(B_{A_1})_{ij} = \mathcal{L}\|\mathbf{x} - \mathbf{x}_{\mathcal{I}_i^1}\|^m|_{\mathbf{x}=\mathbf{x}_{\mathcal{I}_j^1}}, i = 1, \dots, n, j = 1, \dots, p_1, \quad (11)$$

$$(B_{\Psi_1})_{ij} = \mathcal{L}\psi_i^1(\mathbf{x})|_{\mathbf{x}=\mathbf{x}_{\mathcal{I}_j^1}}, i = 1, \dots, M, j = 1, \dots, p_1. \quad (12)$$

W_1 is the matrix of overlapped RBF-FD weights, with each column containing the RBF-FD weights for a point $\mathbf{x} \in \mathbb{B}_1$. It is also notationally useful to refer to the column of W_1 associated with the point \mathbf{y} as $W_1(\mathbf{y})$. The linear system (8) has a unique solution if the nodes in P_1 are distinct and Ψ_1 has full column rank [22, 66]. The matrix of polynomial coefficients W_1^ψ enforces the polynomial reproduction constraint (7). This constraint ensures that the *local* approximation error is bounded by $O(h^{\ell+1-\theta})$, where θ is the order of the differential operator \mathcal{L} , and h is the largest distance between the point at which the weights are computed and every other point in the stencil [16].

2.2. Automatically determining R_1 and \mathbb{B}_1

In previous versions of the overlapped RBF-FD method, the set R_1 and the ball \mathbb{B}_1 were determined by defining a overlap parameter $\delta \in (0, 1]$ such that all nodes within a distance $(1 - \delta)\rho_1$ from the center lay within the ball \mathbb{B}_1 , where $\rho_1 = \max_{1 \leq j \leq n} \|\mathbf{x}_{\mathcal{I}_1^1} - \mathbf{x}_{\mathcal{I}_j^1}\|$. As such an approach requires tuning δ , in this work, we present an automatic approach to determine the sets R_1 and \mathbb{B}_1 . This new approach uses a pair of stability indicators to determine whether a set of computed weights is of sufficient quality. In the discussion that follows, we continue to use \mathbf{y} to refer to the point at which weights $\mathbf{w}^1(\mathbf{y})$ are computed, and \mathbf{x} to refer to points comprising the stencils.

2.2.1. The local \mathcal{L} -Lebesgue function indicator

As described in [49] and noted in [3, 5], local \mathcal{L} -Lebesgue functions play a key role in assessing the suitability of a set of RBF-FD weights. Large values of the local \mathcal{L} -Lebesgue function can lead to spurious eigenvalues in the differentiation matrix corresponding to \mathcal{L} . This fact was used to develop a stability indicator for the overlapped RBF-FD method in [49] which was used to discard unsuitable weights on a given stencil. Recall that $\mathbf{x}_{\mathcal{I}_1^1}$ and its $n - 1$ neighbors form the stencil P_1 . Then, letting $\mathbf{y}_1 = \mathbf{x}_{\mathcal{I}_1^1}$, one can define the \mathcal{L} -Lebesgue function at \mathbf{x}_1 as the ℓ_1 -norm of the weight vector at that point:

$$\Lambda_{\mathcal{L}}(\mathbf{y}) = \|W_1(\mathbf{y})\|_1, \quad (13)$$

where $W_1(\mathbf{y}) = \mathbf{w}^1(\mathbf{y})$ refers to the column of W_1 corresponding to the point \mathbf{y} . We can now define a set \mathbb{B}_{11} as:

$$\mathbb{B}_{11} = \{\mathbf{y} \in P_1 \mid \Lambda_{\mathcal{L}}(\mathbf{y}) \leq \Lambda_{\mathcal{L}}(\mathbf{y}_1)\}, \quad (14)$$

i.e., the set of all points in P_1 where the \mathcal{L} -Lebesgue function takes on values smaller than at the stencil center \mathbf{y}_1 . In [49], the RBF-FD weights from lower-order methods were computed by solving (8) (with the right hand side determined by the overlap parameter), but were tested (and if necessary, discarded) using the \mathcal{L} -Lebesgue stability indicator. In this work, we do not use the overlap parameter, but instead directly compute RBF-FD weights for every $\mathbf{y} \in P_1$, and assess their suitability using the \mathcal{L} -Lebesgue function at that point. This process is repeated on every stencil.

2.2.2. An oscillation indicator

All RBFs have an associated reproducing kernel Hilbert space called the native space [22]. The *native space semi-norm* of an RBF interpolant formed from PHS RBFs is a measure of how much the interpolant oscillates [7]. In the standard interpolation setting (rather than the RBF-FD setting), PHS RBF interpolants with RBF interpolation matrix A and RBF coefficient vector \mathbf{c} have a native space semi-norm of $|\mathbf{c}^T A \mathbf{c}|$ [7]. We now define an analogous oscillation indicator for the RBF-FD context on the stencil P_1 for the interpolant $s_1(\mathbf{x}, \mathbf{y})$ as:

$$\mathcal{S}(\mathbf{y})|_{\mathbf{y} \in P_1} = \left| \begin{bmatrix} W_1(\mathbf{y})^T & (W_1^\psi(\mathbf{y}))^T \end{bmatrix} \begin{bmatrix} A_1 & \Psi_1 \\ \Psi_1^T & O \end{bmatrix} \begin{bmatrix} W_1(\mathbf{y}) \\ W_1^\psi(\mathbf{y}) \end{bmatrix} \right|. \quad (15)$$

This quantity serves as a second stability indicator for the overlapped RBF-FD method. We use it to define a set \mathbb{B}_{12} (analogous to \mathbb{B}_{11}):

$$\mathbb{B}_{12} = \{\mathbf{y} \in P_1 \mid \mathcal{S}(\mathbf{y}) \leq \mathcal{S}(\mathbf{y}_1)\}. \quad (16)$$

While [49] found that the \mathcal{L} -Lebesgue indicator is sufficient for low-order methods, we found that the above oscillation indicator was vital for stability in high-order methods. To better understand this indicator, we can multiply out the matrix and vectors in (15) to obtain:

$$\mathcal{S}(\mathbf{y})|_{\mathbf{y} \in P_1} = \left| W_1(\mathbf{y})^T B_{A_1} + (W_1^\psi(\mathbf{y}))^T B_{\Psi_1} \right|, \quad (17)$$

where B_{A_1} and B_{Ψ_1} are evaluations of \mathcal{L} applied to the basis functions. Since application of $W_1(\mathbf{y})^T$ corresponds to approximating the action of \mathcal{L} , this indicator can be thought of as measuring the magnitude of higher-order derivatives of the basis functions (RBFs and polynomials) as approximated by computed RBF-FD weights. While theoretical justification for this indicator is lacking, our experiments indicate that this indicator works well in conjunction with the \mathcal{L} -Lebesgue indicator.

2.2.3. Parameter-Free Assembly

Once both stability indicators are used to define the sets \mathbb{B}_{11} and \mathbb{B}_{12} , we can now define a single set \mathbb{B}_1 of points whose weights are acceptable from the stencil P_1 as:

$$\mathbb{B}_1 = \mathbb{B}_{11} \cap \mathbb{B}_{12}, \quad (18)$$

i.e., the set of nodes for which we deem the RBF-FD weights as suitable is the intersection of the sets of nodes which pass both stability indicator tests. It is also useful to obtain the global indices of the points each ball \mathbb{B}_1 :

$$R_1 = \{\mathcal{R}_1^1, \mathcal{R}_2^1, \dots, \mathcal{R}_{p_1}^1\}, \quad (19)$$

where $p_1 = |\mathbb{B}_1|$. Note that with the definition of the sets \mathbb{B}_k , the stencil centers $\mathbf{y}_k = \mathbf{x}_{\mathcal{I}_1^k}$ will automatically pass our stability tests, and an assembly algorithm based on the stencil centers is guaranteed to converge in that every point $\mathbf{x} \in X$ is guaranteed to receive a set of RBF-FD weights. The indicators (14) and (16) are both visualized in Figure 1.

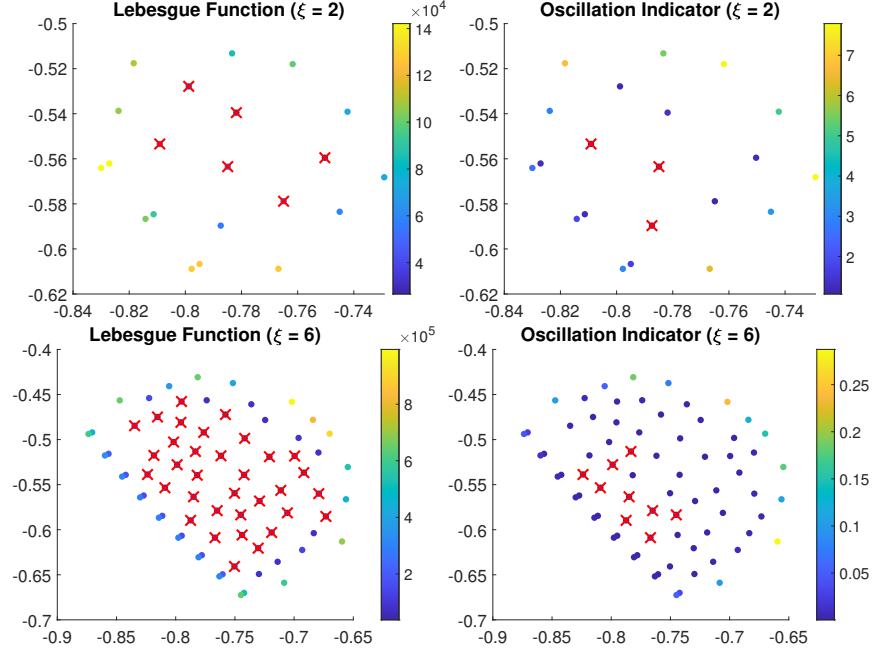


Figure 1: A visualization of (14) (left column) and (16) (right column) for a stencil used to compute RBF-FD weights for the Laplacian on the unit disk ($N = 4975$ nodes) with two embedded ellipses given by (53) and (54). The colors show the function values pointwise, while the crosses indicate the points that have been selected by the indicator. The top row shows both indicators for $\xi = 2$, a second order method, and the bottom row shows the indicators for $\xi = 6$, a sixth-order method.

In practice, even for irregular nodes, we find that the number of stencils N_s can be much smaller than the number of nodes N . In addition, we find that the order θ of the differential operator \mathcal{L} can influence the number of stencils, with higher-order operators leading to fewer stencils. This is likely due to the fact that our formula for the stencil size n depends indirectly on θ , and also due to the fact that larger stencils result in better behaved RBF-FD weights [49]. The full process for assembling the sparse differentiation matrix using the indicators (14) and (16) is described in Algorithm 1. The new algorithm no longer requires an overlap parameter, and is therefore much more robust to irregularities in node sets. More importantly, if the node sets are changing in time (as they do in the present study), this approach obviates the need for hand-tuning.

3. A high-order meshless framework for advection-diffusion equations

We now present a high-order semi-Lagrangian meshless method that takes advantage of the parameter-free overlapped RBF-FD formulation given in Algorithm 1. Our complete method is outlined in Algorithm 2. Algorithm 2 references several other algorithms and sections in this work, which will be explained later in the text. Also, since the overall framework uses a multistep method, it is important to appropriately modify the algorithm for the first two steps; we do this by using single steps of lower-order multistep methods. This section is organized as follows: first, we explain the overarching semi-Lagrangian (SL) ghost node method in Section 3.1. Then, in Section 3.2, we describe the underlying node adaption algorithm used to tackle moving embedded boundaries. Next, in Section 3.3, we explain how to use Algorithm 1 to generate overlapped local RBF interpolation stencils for use within the SL ghost node method. Parameters and error estimates are described in Section 3.5. The preconditioner on step 16 of Algorithm 1 is described in Section 4.1. We defer a full complexity analysis of the algorithm to Section 5.

Algorithm 1 Automatic Differentiation matrix assembly

Given: $X = \{\mathbf{x}_k\}_{k=1}^N$, the set of nodes in the domain.
Given: \mathcal{L} , the linear differential operator to be approximated.
Given: $n \ll N$, the stencil size.
Generate: L , the $N \times N$ differentiation matrix approximating \mathcal{L} on the set X .
Generate: N_s , the number of stencils.

- 1: Build a k-d tree on the set X in $O(N \log N)$ operations.
- 2: Initialize g , an array of N flags set to 0.
- 3: Initialize the stencil counter, $N_s = 0$.
- 4: **for** $k = 1, N$ **do**
- 5: **if** $g(k) == 0$ **then**
- 6: Use k-d tree to determine $\{\mathbf{x}_{\mathcal{I}_1^k}, \dots, \mathbf{x}_{\mathcal{I}_n^k}\}$. Here, $\mathcal{I}_1^k = k$.
- 7: Use (8) to compute W_k , the $n \times n$ matrix of RBF-FD weights on the full stencil P_1 .
- 8: Find the set \mathbb{B}_k using (14) and (16).
- 9: Also find the set R_k by keeping track of global indices of elements of \mathbb{B}_k .
- 10: **for** $i = 1, n$ **do**
- 11: **if** $\mathbf{x}_{\mathcal{R}_i^k} \notin \mathbb{B}_k$ **then**
- 12: CONTINUE.
- 13: **end if**
- 14: Set $g(\mathcal{R}_i^k) = 1$.
- 15: **for** $j = 1, n$ **do**
- 16: Set $L(\mathcal{R}_i^k, \mathcal{I}_j^k) = W_k(j, i)$.
- 17: **end for**
- 18: **end for**
- 19: **end if**
- 20: **end for**

3.1. A ghost node formulation

We now discuss the ghost node formulation used in Algorithm 2. First, we rewrite the advection-diffusion equations in (1) in *Lagrangian form* as

$$\frac{dc}{dt} = \nu \Delta c + f(\mathbf{x}, t), \mathbf{x} \in \Omega(t), \quad (20)$$

where $\frac{d}{dt} = \frac{\partial}{\partial t} + \mathbf{u} \cdot \nabla$ is the *material* or *Lagrangian* derivative. Once the material derivative is discretized in some suitable fashion, the above equation also requires a suitable discretization of the Laplacian Δ and the boundary condition operator $\alpha \mathbf{n} \cdot \nabla + \beta$ from (2) on the time-varying domain $\Omega(t)$. Our approach is to use the backward differentiation formula (BDF) scheme to discretize the above equation in time, with the Laplacian, boundary conditions, and forcing terms treated implicitly. For a stable spatial discretization in the presence of derivative boundary conditions, we use the ghost node scheme outlined in [49, 50]. We now discuss the details of this scheme in the context of domains with time-varying embedded boundaries, though the scheme is easily adapted to domains where the outer boundary also varies in time.

For a given domain $\Omega(t)$, we define a node set $X(t) = \{\mathbf{x}_k(t)\}_{k=1}^{N(t)} \subset \Omega(t)$ that discretizes $\Omega(t)$. This node set is explicitly divided into a set of *interior nodes* $X_i(t)$ and a set of *boundary nodes* $X_b(t)$ with cardinality $N_i(t)$ and $N_b(t)$, respectively. In addition, we tile the boundary nodes some small distance in the outward normal direction to obtain a set of *ghost nodes* $X_g(t)$, also of cardinality $N_b(t)$. This forms an extended node set $X_e(t) = X(t) \cup X_g(t)$, which is used extensively in Algorithm 2.

Our ghost node scheme involves enforcing the PDE up to and including the domain boundary $\partial\Omega(t)$, and enforcing boundary conditions at the boundary. We discretize the Laplacian using Algorithm 1 with $\mathcal{L} \equiv \Delta$ on the extended node set $X_e(t)$. The time-varying discrete Laplacian $L(t)$ can be written in block form as:

$$L(t) = \begin{bmatrix} L_{ii}(t) & L_{ib}(t) & L_{ig}(t) \\ L_{bi}(t) & L_{bb}(t) & L_{bg}(t) \end{bmatrix}, \quad (21)$$

Algorithm 2 Semi-Lagrangian Advection-Diffusion on Moving Domains

Given: X_0 , the initial node set on the interior and boundary of a time-invariant reference domain Ω_0 .

Given: $(X_e)_0 = X_0 \cup (X_g)_0$, the initial extended node set on Ω_0 containing interior, boundary, and ghost nodes.

Given: Seed nodes on all embedded boundaries.

Given: h , the average separation distance between nodes.

Given: ξ , the desired order of approximation of the numerical method.

Given: ν , the diffusion coefficient.

Given: \mathbf{u} , an incompressible velocity field.

Given: $c_0(\mathbf{x}) = c(\mathbf{x}, 0)$, an initial condition.

Given: Δt , the time step.

Given: T , the final time.

Given: $g(\mathbf{x})$, the desired boundary condition such that $\mathcal{B}c = g$ over the full domain boundary.

Generate: $\underline{C} = C(\mathbf{x}, t)|_X \approx c(\mathbf{x}, t)|_X$, the numerical solution to (20).

- 1: Set polynomial degree ℓ , PHS RBF exponent m , and stencil size n according to Table 1.
 - 2: Set $n_s = \left\lfloor \frac{T}{\Delta t} \right\rfloor$ and adjust Δt so that $\Delta t n_s = T$.
 - 3: Reconstruct embedded boundaries from seed nodes (as outlined in Section 3.2).
 - 4: Modify $(X_e)_0$ to account for embedded boundaries to obtain $(X_e)(t_0)$ (as outlined in Section 3.2).
 - 5: Use Algorithm 1 with $\mathcal{L} = \delta \circ$ and $(X)(t_0)$ to obtain (localized) interpolation operator \mathcal{I}^0 (as outlined in Section 3.3).
 - 6: Use Algorithm 1 with $\mathcal{L} = \Delta$ and $\mathcal{L} = \mathcal{B}$ to obtain sparse matrices $L(t_0)$ and $B(t_0)$ respectively.
 - 7: Set $C^0 = c_0(\mathbf{x})$.
 - 8: **for** $k = 1, \dots, n_s$ **do**
 - 9: Set $t_{n+1} = k\Delta t$.
 - 10: Move seed nodes on embedded boundaries using \mathbf{u} and RK3 with the same time-step Δt .
 - 11: Reconstruct embedded boundaries from seed nodes and *modify* $(X_e)_0$ to obtain $(X_e)(t_{n+1})$.
 - 12: *Update* sparse matrices $L(t_n)$ and $B(t_n)$ to $L(t_{n+1})$ and $B(t_{n+1})$ using Algorithm 3.
 - 13: *Update* interpolation operator \mathcal{I}^n to \mathcal{I}^{n+1} using Algorithm 3.
 - 14: Trace back $(X_e)(t_{n+1})$ to t_n, t_{n-1} , and t_{n-2} using (28)–(29) to obtain $(X_e)_d^n, (X_e)_d^{n-1}$, and $(X_e)_d^{n-2}$.
 - 15: Compute $(C)_d^n = \mathcal{I}^n C^n$, $(C)_d^{n-1} = \mathcal{I}^{n-1} C^{n-1}$, and $(C)_d^{n-2} = \mathcal{I}^{n-2} C^{n-2}$ using the procedure outlined in Section 3.3.
 - 16: Form the preconditioner as outlined in Section 4.1.
 - 17: Form and solve the BDF3 linear system (26) (or its BDF1 or BDF2 analogues if $k = 1$ or $k = 2$) to obtain C^{n+1} .
 - 18: Set $C^{n-2} = C^{n-1}$, $C^{n-1} = C^n$, and $C^n = C^{n+1}$.
 - 19: Set $\mathcal{I}^{n-2} = \mathcal{I}^{n-1}$, $\mathcal{I}^{n-1} = \mathcal{I}^n$, and $\mathcal{I}^n = \mathcal{I}^{n+1}$.
 - 20: **end for**
-

where the subscripts indicate partitions of the Laplacian corresponding to interior (i), boundary (b), and ghost (g) points. Notice that $L(t)$ is computed only at the interior and boundary points, but uses stencils that involve ghost points, giving the matrix dimensions of $N(t) \times (N(t) + N_b(t))$, where $N(t) = N_i(t) + N_b(t)$. We also discretize the boundary condition operator using Algorithm 1 with $\mathcal{L} = \alpha(\mathbf{x}, t)\mathbf{n} \cdot \nabla + \beta(\mathbf{x}, t)$. The resulting discrete boundary operator $B(t)$ can be written as:

$$B(t) = [B_{bi}(t) \quad B_{bb}(t) \quad B_{bg}(t)]. \quad (22)$$

This matrix has dimensions $N_b(t) \times (N(t) + N_b(t))$. The sparse matrices $L(t)$ and $B(t)$ can now be used to discretize the advection-diffusion equation.

Let $C(\mathbf{x}, t) \approx c(\mathbf{x}, t)$ be the numerical solution to the advection-diffusion equation. First, partition $C(\mathbf{x}, t)|_{X_e}$ into $C_i = C(\mathbf{x}, t)|_{X_i}$, $C_b = C(\mathbf{x}, t)|_{X_b}$, and $C_g = C(\mathbf{x}, t)|_{X_g}$. We can use these partitions to write the

discretized advection-diffusion equation as:

$$\frac{dC_i}{dt} = \nu (L_{ii}(t)C_i + L_{ib}(t)C_b + L_{ig}(t)C_g) + f_i(t), \quad (23)$$

$$\frac{dC_b}{dt} = \nu (L_{bi}(t)C_i + L_{bb}(t)C_b + L_{bg}(t)C_g) + f_b(t), \quad (24)$$

$$B_{bi}(t)C_i + B_{bb}(t)C_b + B_{bg}(t)C_g = g_b(t), \quad (25)$$

where $g_b(t) = g(\mathbf{x}, t)|_{X_b(t)}$. The above system can only be treated as a set of ODEs if a suitable discretization for the material derivative $\frac{d}{dt}$ is used. For example, for the third-order BDF scheme (BDF3) [1], we obtain:

$$\underbrace{\begin{bmatrix} I - \frac{6}{11}\nu\Delta t L_{ii}^{n+1} & -\frac{6}{11}\nu\Delta t L_{ib}^{n+1} & -\frac{6}{11}\nu\Delta t L_{ig}^{n+1} \\ -\frac{6}{11}\nu\Delta t L_{bi}^{n+1} & I - \frac{6}{11}\nu\Delta t L_{bb}^{n+1} & -\frac{6}{11}\nu\Delta t L_{bg}^{n+1} \\ B_{bi}^{n+1} & B_{bb}^{n+1} & B_{bg}^{n+1} \end{bmatrix}}_{A(t_{n+1})} \underbrace{\begin{bmatrix} C_i^{n+1} \\ C_b^{n+1} \\ C_g^{n+1} \end{bmatrix}}_{C^{n+1}} = \underbrace{\begin{bmatrix} \frac{18}{11}(C_i)_d^n - \frac{9}{11}(C_i)_d^{n-1} + \frac{2}{11}(C_i)_d^{n-2} + \Delta t \frac{6}{11}f_i^{n+1} \\ \frac{18}{11}(C_b)_d^n - \frac{9}{11}(C_b)_d^{n-1} + \frac{2}{11}(C_b)_d^{n-2} + \Delta t \frac{6}{11}f_b^{n+1} \\ g_b^{n+1} \end{bmatrix}}_{r^{n+1}}, \quad (26)$$

where the superscripts now indicate time levels. The subscript d under a variable denotes the value of that variable at an SL departure point [68]. We explain this in greater detail, focusing without loss of generality on the interior points. Recall that $C_i^n = C(\mathbf{x}, t_n)|_{X_i(t_n)}$, where $X_i(t_n)$ is the set of interior points at time level n . Then, the variable $(C_i)_d^n$ can be written as:

$$(C_i)_d^n = C(\mathbf{x}, t_n)|_{(X_i)_d(t_n)}, \quad (27)$$

where the set of *departure points* $(X_i)_d(t_n) = \{(\mathbf{x}_j)_d(t_n)\}_{j=1}^{N_i(t_n)}$ is defined by solving the following set of ODEs *backward* in time:

$$\frac{d\mathbf{p}_j}{dt} = \mathbf{u}(\mathbf{p}_j, t), \quad (28)$$

$$\mathbf{p}_j(t_{n+1}) = \mathbf{x}_j(t_{n+1}), j = 1, \dots, N_i(t_{n+1}). \quad (29)$$

These ODEs can be solved using a standard numerical ODE solver; we use the third-order Runge-Kutta (RK3) method. This process is called *trajectory reconstruction* [53], since it reconstructs the trajectory a (fictitious) particle would take if it arrived at the nodes $X_i(t_{n+1})$; it is also alternatively referred to as the back trace procedure. In simple terms, solving the above set of ODEs for each of the $N_i(t_{n+1})$ nodes results in the set of departure points $(X_i)_d(t_n)$. In general, however, the set of departure points $(X_i)_d(t_n)$ differs from the set of interior nodes $X_i(t_n)$ at time level n . Consequently, C_i^n must be interpolated to $(X_i)_d(t_n)$ to obtain $(C_i)_d^n$. For convenience, let us define an abstract time-dependent interpolation operator $\mathcal{I}_i(t, \cdot, \cdot)$ such that

$$(C_i)_d^n = \mathcal{I}_i(t_n, C_i^n, C_b^n), \quad (30)$$

where $\mathcal{I}_i(t, \cdot, \cdot)$ is an interpolation operator that lets us interpolate fields from $X_i(t_n)$ and $X_b(t_n)$ to $(X_i)_d(t_n)$. We defer discussion of this operator to Section 3.3. To simplify the notation, we set

$$\mathcal{I}_i(t_n, C_i^n, C_b^n) = \mathcal{I}_i^n(C_i^n, C_b^n).$$

Note that Algorithm 2 also requires us to compute the quantities $(C_i)_d^{n-1}$ and $(C_i)_d^{n-2}$. In analogy with (30), these quantities can be written as:

$$(C_i)_d^{n-1} = C(\mathbf{x}, t_{n-1})|_{(X_i)_d(t_{n-1})} = \mathcal{I}_i^{n-1}(C_i^{n-1}, C_b^{n-1}), \quad (31)$$

$$(C_i)_d^{n-2} = C(\mathbf{x}, t_{n-2})|_{(X_i)_d(t_{n-2})} = \mathcal{I}_i^{n-2}(C_i^{n-2}, C_b^{n-2}), \quad (32)$$

where the operators \mathcal{I}_i^{n-1} and \mathcal{I}_i^{n-2} interpolate quantities from X_i and X_b to $(X_i)_d$ at time levels t_{n-1} and t_{n-2} respectively. To find the departure points $(X_i)_d(t_{n-1})$ and $(X_i)_d(t_{n-2})$, we solve (28) backward

in time to levels $n - 1$ and $n - 2$. This requires advecting/tracing nodes $(X_i)(t_{n+1})$ backward for several steps [68]. The same approach can be used to obtain $(C_b)_d$ values at different time levels by defining boundary interpolation operators $\mathcal{I}_b(t)$.

Since the solution of the diffusion problem is done on the node set $X(t_{n+1})$, this approach proves far more convenient than purely Eulerian methods, which typically require spatial extrapolation to fill points that entered the domain in the current step (*e.g.*, see [41]); this issue is completely avoided in the SL approach.

Remark 1. In Algorithm 2, we move the seed nodes on the embedded boundaries using an RK3 discretization of (28) with the seed node positions in place of the nodes \boldsymbol{x} , but *forward* in time. In practical scenarios, it is straightforward to replace this step with updates of the type seen in the IB method [56].

In the following subsections, we discuss how to generate the node sets $X_e(t)$, construct and use the interpolation operators $\mathcal{I}(t)$, and efficiently *update* both the discrete Laplacians $L(t)$ and the interpolation operators $\mathcal{I}(t)$.

3.2. Node set adaptation

Our technique for solving the advection-diffusion equation on a time-varying domain $\Omega(t)$ involves using Algorithm 1 to assemble the discrete Laplacian $L(t)$ and the discrete boundary operator $B(t)$ on a *time-varying node set* $X_e(t)$. In this work, we use the node generation and adaptation algorithm described in [51], adapted and optimized for moving embedded boundaries. We make this choice because the algorithm from [51] is designed to *locally* adapt nodes around embedded boundaries, allowing us to reuse previously computed interpolation operators and overlapped RBF-FD weights at nodes that are sufficiently far away from the moving embedded boundaries; this is discussed in greater detail in Section 3.4. We assume in the following discussion that the (irregular) outer boundary stays fixed over time, though our algorithm can be modified trivially to handle the case where the outer boundary also moves.

Let $\Omega(t)$ be a time-varying domain defined using a time-invariant reference domain Ω_0 and N_Γ time-varying subdomains $\{\Omega_j(t)\}_{j=1}^{N_\Gamma}$ so that:

$$\Omega(t) = \Omega_0 \setminus \bigcup_{j=1}^{N_\Gamma} \Omega_j(t). \quad (33)$$

The domain boundary $\Gamma(t)$ can then be written as:

$$\Gamma(t) = \bigcup_{j=0}^{N_\Gamma} \Gamma_j(t), \quad (34)$$

where $\Gamma_j(t), j = 1, \dots, N_\Gamma$ are the boundaries of the time-varying subdomains $\Omega_j(t)$, and $\Gamma_0(t) = \Gamma_0$ is fixed for all time. In this setting, our goal is now to generate the set $X_e(t)$, which involves generating interior nodes $X_i(t)$, boundary nodes $X_b(t)$, and ghost nodes $X_g(t)$. Our approach is to generate node sets on Ω_0 and Γ_0 prior to starting to advance the time-dependent solution, and then adapt these to account for the time-varying subdomains $\Omega_1, \dots, \Omega_{N_\Gamma}$ and their boundaries $\Gamma_1, \dots, \Gamma_{N_\Gamma}$.¹ The approach is as follows:

1. Starting from a small set of seed nodes on Γ_0 , generate the node set $X_0 = (X_i)_0 \cup (X_b)_0$ for $\Omega_0 \cup \Gamma_0$ and a geometric representation for Γ_0 using Algorithm 1 from [51]. Use this geometric representation to generate a set of outward unit normal vectors $\mathcal{N}_0 = \{(\boldsymbol{n}_0)_j\}_{j=1}^{N_b}$ and use them, in turn, to generate the ghost nodes $(X_g)_0$ for the domain Ω_0 . This gives us the extended node set $(X_e)_0$ on the time-invariant domain $\Omega_0 \cup \Gamma_0$.

¹If instead a node set is given directly on $\Omega(t_0)$, it is sufficient for our methods to precede the node adaptation procedure with a step that simply fills $\Omega_j(t_0), j = 1, \dots, N_\Gamma$ with nodes at the initial time $t = t_0$.

2. Next, at any time t , use Algorithm 5 from [51] to adapt the node set by “turning-off” any nodes contained in $\bigcup_{j=1}^{N_\Gamma} \Omega_j(t)$. This requires forming a geometric representation of the boundaries $\bigcup_{j=1}^{N_\Gamma} \Gamma_j(t)$ from a set of initially quasi-uniform seed nodes spaced h_d apart. We use the updated positions of the seed nodes and the geometric modeling technique presented in [51] to form geometric representations of the embedded boundaries. This node adaptation algorithm uses normal vectors on each subdomain boundary to test whether the nodes in $(X_i)_0$ are inside or outside the domain $\Omega(t) = \Omega_0 \setminus \bigcup_{j=1}^{N_\Gamma} \Omega_j(t)$.

The geometric modeling error is $O(h_d^8)$ [51].

3. We now need to ensure that the boundary nodes $X_b(t)$ respect the average node spacing h . While the outer boundary nodes are assumed to be fixed in time, the inner boundaries Γ_j may deform over the course of a simulation. In Algorithm 2, we describe moving the boundaries by moving the seed nodes; this amounts to a Lagrangian description of the embedded boundaries, not unlike in the IB method [45]. To avoid boundary nodes being further apart than h due to this movement, we first reconstruct the embedded boundaries from their seed nodes using the previously-mentioned geometric modeling technique. We then use Algorithm 2 from [51] to sample the embedded boundaries in such a way that the resulting node sets are quasi-uniform and neighboring points a distance of approximately h apart. Thus, in Algorithm 2 (as used in this article), the seed nodes on the embedded boundaries are specified once at the beginning of the simulation (or at the introduction of an embedded boundary). In contrast, the boundary nodes are **regenerated** every time step and may vary in number. If the seed nodes drift apart, one could use the boundary representation at a given step to regenerate the seed nodes via quasi-uniform sampling with a spacing of h_d . Once the the boundary nodes are obtained, the normal vectors at the boundary nodes $X_b(t)$ can be used to obtain the set of ghost nodes $X_g(t)$ (by extension in the normal direction). This generates the fully extended set $X_e(t)$.

3.3. Overlapped local RBF interpolation

During every step of the method, the departure points $(X_i)_d(t_n)$, $(X_i)_d(t_{n-1})$, and $(X_i)_d(t_{n-2})$ (and their boundary counterparts) are calculated by solving (28). Once these departure points are calculated, the solution is interpolated to these departure points using the interpolation operators \mathcal{I}^n , \mathcal{I}^{n-1} , and \mathcal{I}^{n-2} . In this article, we use an analogue of RBF-FD (local RBF interpolation) to compute these interpolation operators, making them completely localized and inexpensive to compute and evaluate. In previous work on a static domain (the sphere), the authors used a separate stencil for each of the points in the domain [53]. Our approach here is similar, except that we now use overlapped RBF-FD to generate potentially fewer stencils.

At a high level, our approach is to decompose the points $X(t) \subset \Omega(t)$ into stencils first, use the back trace procedure to determine which stencil a departure point lies on, then interpolate the solution from the node set $X(t)$ to the set of departure points $X_d(t) = (X_i)_d(t) \cup (X_b)_d(t)$. To automatically determine stencils, we use Algorithm 1 with the point evaluation operator δ_\circ in place of \mathcal{L} . However, unlike in the case of the matrices $L(t)$ and $B(t)$, we do not assemble the resulting weights into a sparse matrix. Instead, we only store the LU decompositions of the interpolation matrices. Once the stencils are determined and the LU decompositions are stored, we build an acceleration structure (a k-d tree) on the stencil centers.

For each time-step, we then determine which stencil each departure point $(\mathbf{x}_j)_d$, $j = 1, \dots, N(t)$ lies on by finding the closest stencil center to that departure point. Suppose that some number of departure points $(\mathbf{x}_j)_d$, $j = 1, \dots, N_k$ turn out to be associated with the stencil P_k , and for simplicity, ignore time levels for the moment. Our goal is to find $C((\mathbf{x}_j)_d, t)$, $j = 1, \dots, N_k$ given the values of $C(\mathbf{x}, t)$ on the stencil P_k and the LU decomposition of the interpolation matrix on P_k . The procedure to do so is as follows:

1. Using the stored LU decomposition of that stencil interpolation matrix from (8) and $C(\mathbf{x}, t)|_{P_k}$ as the right hand side, solve a local linear system for $O(n^2)$ operations to determine a set of RBF and polynomial coefficients \mathbf{c}^k on the stencil P_k .
2. Using these coefficients, evaluate the local interpolant on the stencil P_k at each departure point $(\mathbf{x}_j)_d$, $j = 1, \dots, N_k$ on P_k to get $C((\mathbf{x}_j)_d, t)$.

This procedure is repeated for every stencil.

Finally, it is important to note that we do not use the extended node set $X_e(t)$ to compute the interpolation stencils using Algorithm 1. This would produce interpolation stencils that use unphysical values at ghost nodes to interpolate data to the departure points. Instead, the interpolation stencils are only calculated on the set $X(t) = X_i(t) \cup X_b(t)$, which does not contain ghost nodes. While this results in one-sided stencils at the domain boundary, we found this to be more stable over a wide range of Peclet numbers and boundary conditions than an approach that allowed the incorporation of ghost nodes into SL interpolation stencils.

3.4. Selective updates to RBF-FD weights and stencils

Algorithm 3 Efficient differentiation matrix update

- Given:** $X_e(t_n)$, the $N_e(t) \times d$ matrix of nodes at time t_n .
Given: $X_e(t_{n+1})$, the $N_e(t_{n+1}) \times d$ matrix of nodes at time t_n .
Given: \mathcal{H}_e , a kd-tree built on $X_e(t_{n+1})$.
Given: \mathcal{L} , the linear differential operator to be approximated.
Given: $n \ll N$, the stencil size.
Given: $L(t_n)$, the $N_e(t_n) \times N_e(t_n)$ differentiation matrix corresponding to \mathcal{L} at time t_n .
Given: $X_c(t_n)$, the $N_c \times d$ matrix of *stencil centers* at time t_n .
Given: $N_c(t_n) = |X_c(t_n)|$, the number of stencil centers at time t_n .
Given: \mathcal{N}_c , an $N_c(t_n) \times n$ matrix of nearest neighbor indices for $X_c(t_n)$ in $X_e(t_n)$.
Given: \mathcal{W}_c , a matrix mapping each row of $X_c(t_n)$ to the corresponding rows and columns in $L(t_n)$.
Generate: $L(t_{n+1})$, the $N_e(t_{n+1}) \times N_e(t_{n+1})$ differentiation matrix approximating \mathcal{L} at time t_{n+1} .
- 1: Build a kd-tree \mathcal{H}_c on the set $X_c(t_n)$ in $O(N_c \log N_c)$ operations.
 - 2: Initialize $X_a = X_e(t_{n+1})$, an $N_e(t_n) \times d$ matrix of *active points* (points for which weights must be computed).
 - 3: **for** $k = 1, N_c$ **do**
 - 4: Use \mathcal{H}_e to determine if $X_c(k, :) \in X_e(t_{n+1})$. If not, CONTINUE.
 - 5: Use \mathcal{H}_e to determine n nearest neighbors of $X_c(k, :)$ in $X_e(t_{n+1})$.
 - 6: Query $\mathcal{N}_c(k, :)$ to obtain n nearest neighbors of $X_c(k, :)$ at time t_n .
 - 7: **if** Neighbors at t_{n+1} do not match neighbors at t_n **then**
 - 8: CONTINUE.
 - 9: **end if**
 - 10: Consult $\mathcal{W}_c(k)$ to obtain p_k row indices (stored in $\mathcal{R}_k(t_n)$) and n column indices into $L(t_n)$ (stored in $\mathcal{I}_k(t_n)$) corresponding to $X_c(k, :)$.
 - 11: Using \mathcal{H}_e , get new p_k row indices (stored in $\mathcal{R}_k(t_{n+1})$) and n column indices (stored in $\mathcal{I}_k(t_{n+1})$) into $L(t_{n+1})$.
 - 12: Set $L_{\mathcal{R}_k(t_{n+1}), \mathcal{I}_k(t_{n+1})}(t_{n+1}) = L_{\mathcal{R}_k(t_n), \mathcal{I}_k(t_n)}(t_n)$.
 - 13: Remove $(X_e)_{\mathcal{R}_k(t_{n+1})}(t_{n+1})$ from the active point matrix X_a .
 - 14: **end for**
 - 15: For all remaining points in X_a , use Algorithm 1 to compute the remaining rows of $L(t_{n+1})$.
-

We now discuss our technique for efficiently computing the matrices $L(t)$ and $B(t)$ (in (23)–(25)), and the operators $\mathcal{I}(t)$ on the time-varying node set $X_e(t)$. Since the node sets vary in time, the number of rows and columns and also the entries of $L(t)$ and $B(t)$ are time-varying. It is instructive to enumerate all the scenarios in which these changes occur. For the following, assume that $X_c(t)$ is the set of stencil centers, where a “stencil center” is always a node from the set $X(t)$.² Note that any of the following scenarios could require recomputation of weights for multiple rows of the matrix $L(t)$ since overlapped RBF-FD uses the same stencil to compute weights for multiple rows at once.

²In the standard RBF-FD method, every node from the set $X(t)$ is a stencil center. Since overlapped RBF-FD is used, there are far fewer stencil centers than nodes.

1. A node that was a stencil center at time t_n may leave the domain at time t_{n+1} by becoming covered by one of the embedded domains $\Omega_j(t)$. The set of weights associated with that stencil is no longer valid.
2. A node that was a nearest neighbor to a stencil center at time t_n may no longer be a nearest neighbor to that same stencil center at time t_{n+1} . This could occur either if the node left the domain or another node (such as a boundary point) ended up closer. Again, the corresponding overlapped RBF-FD weights need to be recomputed.
3. The global indices of each stencil's nodes into the set $X_e(t)$ could change since the total number of nodes $N(t) + N_b(t)$ is a function of time. In this case, new global indices must be found so that the previously computed weights are copied in the appropriate rows and columns of $L(t)$.
4. New nodes may be introduced to the domain. For instance, nodes on the moving boundaries $\Gamma(t)$ are always changing, as are their ghost nodes. All these new nodes create rows in $L(t)$. Overlapped RBF-FD weights must be freshly computed for each of these nodes.
5. In the case of $B(t)$, the differentiation matrix that enforces boundary conditions, we must therefore not only check the above cases, but also check if α and/or β have changed. If they have, we need to recompute the RBF-FD weights for the boundary.

All of the above statements for $L(t)$ also apply to $\mathcal{S}(t)$ with a minor caveat.³ Algorithm 3 shows this procedure for the Laplacian, but can be trivially adapted to update the boundary condition matrix and interpolation operators also. All the nodes for which weights are not copied are marked as such, and Algorithm 1 is then applied to compute the weights for these nodes.

Remark 2. While Algorithm 3 describes the update procedures in terms of matrices, it is again easily adapted to a matrix-free approach where the matrices $L(t)$ and $B(t)$ are never formed, but are instead simply *applied* to approximate solution vectors.

3.5. Error Estimates

We now discuss the error estimates for the discrete PDE (26). To the best of our knowledge, error estimates for SL methods are only available for pure linear advection problems [20]. A formal analysis in the context of a multistep method is beyond the scope of this article. However, it is possible to heuristically account for the different sources of error in (26) and combine them to get an estimate of the total error.

First, we account for all spatial errors due to RBF interpolation or RBF-FD approximations. Recall that h is a measure of average node spacing. For the overlapped RBF-FD method (including overlapped local RBF interpolation), the spatial error in approximating a differential operator of order θ is $O(h^{\ell_{\text{op}}+1-\theta})$, where ℓ_{op} is the degree of the polynomial used within the RBF-FD (or local RBF) formula for that operator [16]. In our case, $\theta = 2$ for the Laplacian, $\theta = 1$ for derivative boundary conditions (Neumann or Robin), and $\theta = 0$ for interpolation. Let E_{Δ} be the error in approximating the Laplacian, and E_{∇} the error in approximating the gradient. These can now be written as:

$$E_{\Delta} = O(h^{\ell_{\Delta}-1}), \quad E_{\nabla} = O(h^{\ell_{\nabla}}), \quad (35)$$

where ℓ_{Δ} and ℓ_{∇} are the polynomial degrees used for computing the overlapped RBF-FD weights for the operators Δ and ∇ respectively.

Next, we account for the errors from the SL portion of the algorithm. (26) relies on three back traces and interpolations per step. Let ℓ_I be the degree of the polynomial used for overlapped local RBF interpolation. Then, the total error from the SL portion of the algorithm is given by [20]:

$$E_{SL} = O\left(\Delta t^p + \frac{h^{\ell_I+1}}{\Delta t}\right). \quad (36)$$

³The only caveat is that no large sparse matrix is maintained for $\mathcal{S}(t)$, and only the stencil information and decomposed local interpolation matrices are needed.

where p is the order of the method used to perform the backtrace ($p = 3$ for RK3). Finally, we have an error contribution of $O(\Delta t^3)$ from the BDF3 time-stepping scheme itself. The total global error for stepping the PDE on the interior and the boundary can be written as:

$$E_{Tot} = O(h^{\ell_\Delta-1}) + O(h^{\ell_\nabla}) + O(\Delta t^3) + O\left(\Delta t^p + \frac{h^{\ell_I+1}}{\Delta t}\right). \quad (37)$$

Let ξ be the desired spatial approximation order of the method. By setting $\ell_\Delta = \xi + 1$ and $\ell_\nabla = \xi$, *i.e.*, choosing different polynomial degrees for each differential operator, the first two terms become h^ξ . In the last term, we set $\Delta t = O(h)$, set $\ell_I = \xi$, and $p = 3$. The error estimate therefore simplifies to

$$E_{Tot} = O(h^\xi) + O(\Delta t^3). \quad (38)$$

These parameter choices for Algorithm 2 are summarized in Table 1.

Remark 3. Our choice of $\Delta t = O(h)$ was *not* motivated by stability, but rather by the desire to have low temporal error.

Remark 4. While we do not have a proof of stability, it is well known that the SL framework allows for time-steps that are much larger than the CFL limit, including in the RBF context [53]. However, in order to prevent trajectories from crossing, it is common to set $\Delta t \leq |J|^{-1}$, where $|J|$ is the minimum pointwise Jacobian of the velocity field evaluated on the collocation node set. In this work, we find that setting $\Delta t = \frac{0.3h}{U_{max}}$ is sufficient for both stability and accuracy, where U_{max} is an estimate of the spatiotemporal maximum of $\|\mathbf{u}\|_2$.

Remark 5. The above analysis ignores geometric modeling errors due to moving boundaries, which are $O(h_d^8)$. These errors are significantly smaller than the other errors in our work and can be safely ignored.

Parameter	Meaning	Value
ℓ_I	Polynomial degree for Interpolation	ξ
ℓ_∇	Polynomial degree for Neumann/Robin BCs	ξ
ℓ_Δ	Polynomial degree for Laplacian	$\xi + 1$
m_{op}	PHS degree for operator op	ℓ_{op} if ℓ_{op} is odd, $\ell_{op} - 1$ if ℓ_{op} is even, $m_{op} = \max(m_{op}, 3)$.
n_{op}	Stencil size for operator op	$2\binom{\ell_{op}+d}{d} + 1$

Table 1: Table of parameters based on desired approximation order ξ , dimension d , and operator $op = \Delta, \nabla$, or I is the operator being approximated.

4. Iterative methods for the implicit system

We now discuss the iterative method we use for solving the system (26). The block matrix $A(t)$ is sparse with at most n non-zero entries per row (where n is the stencil size), and changes every time-step. While the update schemes outlined in Section 3.4 enable fast computation of the matrices $L(t)$ and $B(t)$ that make up $A(t)$, and the right-hand-side of (26), the solution of this time-varying linear system requires both efficient solvers and preconditioners, especially for problems in 3D.

4.1. An efficient saddle-point preconditioner

We use the Generalized Minimum Residuals (GMRES) method [48] to solve (26). However, it is well known that this method requires a good preconditioner to achieve faster convergence to a tolerance. We use a technique outlined in [9], and discuss it briefly here. For what follows, it is useful to write $A(t)$ as a 2×2 block matrix of the form:

$$A(t) = \begin{bmatrix} A_{11}(t) & A_{12}(t) \\ A_{21}(t) & A_{22}(t) \end{bmatrix}, \quad (39)$$

where the block $A_{11}(t)$ is given by

$$A_{11}(t_{n+1}) = \begin{bmatrix} I - \frac{6}{11}\nu\Delta t L_{ii}^{n+1} & -\frac{6}{11}\nu\Delta t L_{ib}^{n+1} \\ -\frac{6}{11}\nu\Delta t L_{bi}^{n+1} & I - \frac{6}{11}\nu\Delta t L_{bb}^{n+1} \end{bmatrix}, \quad (40)$$

and the other blocks naturally follow. The matrix $A(t)$ is a (generalized) saddle-point matrix, with each of its blocks being a sparse matrix, and its *Schur complement* is

$$S(t) = A_{22}(t) - A_{21}(t)A_{11}(t)^{-1}A_{12}(t). \quad (41)$$

$A(t)$ has an inverse iff $S(t)$ is invertible, which in turn requires that $A_{11}(t)$ is invertible [8]. While this appears to be the case in practice for the overlapped RBF-FD method, we now use $S(t)$ to develop a diagonal preconditioner. Consider the block diagonal matrix $\tilde{P}(t)$ given by:

$$\tilde{P}(t) = \begin{bmatrix} A_{11}(t) & O \\ O & S(t) \end{bmatrix}. \quad (42)$$

As described in [9], the *inverse* of \tilde{P} can be a reasonable preconditioner for a system involving $A(t)$. However, computing the true inverse of $\tilde{P}(t)$ requires inverting both $A_{11}(t)$ and $S(t)$. To make this process efficient, we replace $A_{11}(t)$ and $S(t)$ with diagonal matrices. First, define the diagonal matrix $\tilde{A}_{11}(t) = (A_{11})_{ii}$, $i = 1, \dots, N(t)$. Next, define the approximate Schur complement $\tilde{S}(t)$ as $\tilde{S}(t) = A_{22}(t) - A_{21}(t)\tilde{A}_{11}(t)^{-1}A_{12}(t)$. The preconditioner we use is then given by the *inverse* of the diagonal matrix

$$P(t) = \begin{bmatrix} \tilde{A}_{11} & O \\ O & \hat{S}(t) \end{bmatrix}, \quad (43)$$

where $(\hat{S}(t))_{ii} = (\tilde{S}(t))_{ii}$, $i = 1, \dots, N_b(t)$. The inverse of $P(t)$ can be computed efficiently while still serving as a reasonable preconditioner to $A(t)$. This efficiency is important, since $A(t)$ (and hence $P(t)$) changes size every time-step. In practice, we first equilibrate $A(t)$ using Matlab's built-in equilibrate function [18] to permute and rescale A so that its off-diagonal entries are not greater than 1 in magnitude and its diagonal entries are only 1 or -1. We then form $P(t)$ using the blocks of the resulting permuted matrix $B(t)$. The resulting spectra of B and PB are shown in Figure 2 for $\xi = 2$ and $\xi = 6$ (with $\nu = 1$). Our preconditioner has the effect of shifting most of the eigenvalues of B to one side of the imaginary axis. We found in practice that this decreased the number of GMRES iterations by half. This preconditioner is clearly not optimal, and can be improved by adding more blocks from the original matrix to $P(t)$. We leave a deeper investigation of preconditioners for future work.

4.2. A good guess for GMRES

For GMRES to converge rapidly in solving (26), it is also important to supply a good initial guess to the solver. In previous work, the authors have used the solution C^n from the previous time level to good effect [49, 50]. However, for a time-varying domain $\Omega(t)$, the solution C^n lies on the domain $\Omega(t_n)$, while the solution C^{n+1} lies on $\Omega(t_{n+1})$. Consequently, the lengths of the vectors $C(\mathbf{x}, t_n)|_{\mathbf{x} \in X_e(t_n)}$ and $C(\mathbf{x}, t_{n+1})|_{\mathbf{x} \in X_e(t_{n+1})}$ are different, since the node sets $X_e(t) \subseteq \Omega(t)$ themselves vary in time. On the other hand, we see that the vectors $(C_i)_d^n$ and $(C_b)_d^n$ generated from the SL trajectory reconstruction possess cardinality $N_i(t_{n+1})$ and $N_b(t_{n+1})$, respectively. We therefore use these as initial guesses for C_i^{n+1} and C_b^{n+1} , respectively. To obtain an initial guess for C_g^{n+1} , we approximate the value of C_g at the ghost node departure points obtain by back tracing from the ghost nodes at t_{n+1} , *i.e.*, we compute $(C_g)_d^n \approx C(\mathbf{x}, t)|_{(X_g)_d(t_n)}$. In practice, we compute this quantity by evaluating the local overlapped RBF interpolants from Section 3.3 at the ghost departure points. We found that use of this guess vector accelerated GMRES when compared to using the zero vector as a guess (the number of iterations were decreased by an order of magnitude in 3D).

5. Complexity Analysis

We now analyze the computational complexity of Algorithm 2 in terms of (1) the preprocessing (steps 1-7), and (2) the actual time-stepping loop (steps 8-20). Steps 1, 2, and 7 merely involve function evaluations and will be ignored for the purposes of simplicity. For the remainder of the section, we assume without loss of generality that the stencil size n and the polynomial degree ℓ are the same for all operators. All our estimates nevertheless hold true in the worst case sense.

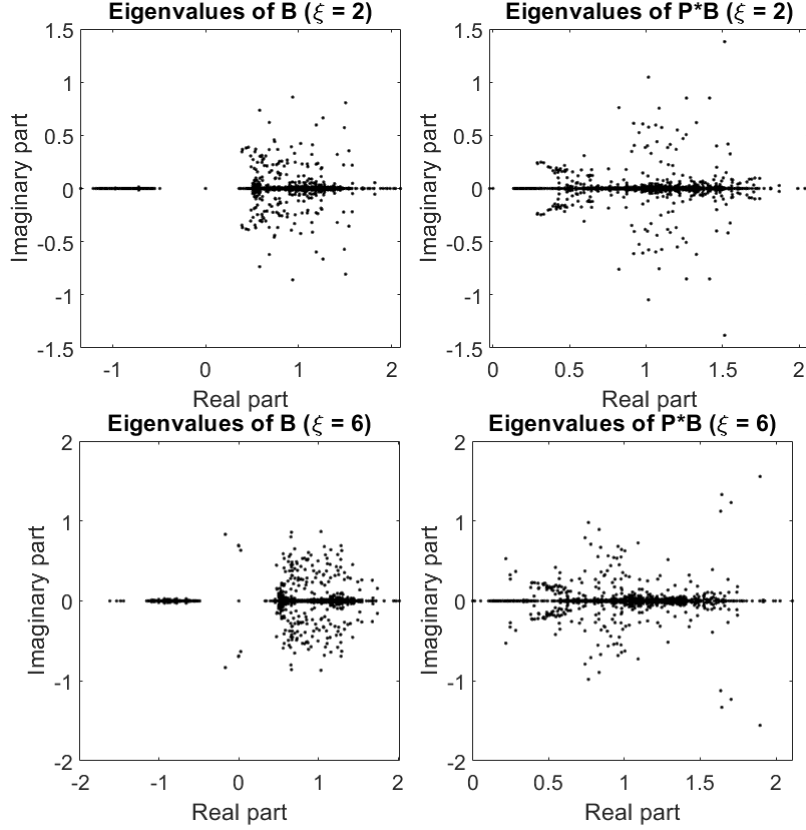


Figure 2: Effect of the preconditioner P from (43) on the equilibrated matrix B for $\xi = 2$ (top row) and $\xi = 6$ (bottom row) for a 2D advection-diffusion problem with $\nu = 1$ at time $t = 4 \times 10^{-6}$.

5.1. Preprocessing complexity

Consider first step 3, which involves interpolation to form the geometric model of the embedded boundaries, and step 4 which involves removal of nodes from the set $(X_e)_0$. Recall that we have N_Γ boundaries, and assume that there are N_d seed nodes on each boundary. Since the geometric model from [51] involves a dense matrix solve, the total cost of *forming* this geometric model is $O(N_\Gamma N_d^3)$. In addition, step 4 involves both evaluation of the geometric model to generate boundary points and a subsequent modification of $(X_e)_0$. Let $N_0 = |(X_e)_0|$. Following Section 4.2 of [51], the costs of steps 3 and 4 can be rewritten as $O(N_\Gamma N_0)$, *i.e.*, linear in N_0 (this relationship can be derived by rewriting the N_d^3 term in terms of N_0 [51]). This is applicable in both 2 and 3 dimensions.

The complexity analysis of steps 5 and 6 depends on the behavior of Algorithm 1, which in turn depends on the type of node set and the differential operator being approximated. Given the polynomial degree ℓ , the number of polynomial terms $M = \binom{\ell+d}{d}$, and the stencil size $n = 2M + 1$, Algorithm 1 will produce a worst case complexity of $O((n + (n - 1)/2)^3 N_0)$, where the cubic term comes from the LU decomposition of (8) and N_0 is the number of points on the reference domain. However, in practice, far fewer linear systems than N_0 will be solved, especially as n and ℓ are increased, since each stencil will be used to compute weights for more than one of the N_0 nodes. To estimate this cost, assume that for each n -node stencil, Algorithm 1 computes RBF-FD weights for $\frac{n}{\kappa}$ nodes, where $1 \leq \kappa \leq n$. Letting N_s be the number of stencils generated by Algorithm 1, the total cost of Algorithm 1 can be written as

$$T_{assembly} = O\left(N_s \left((n + (n - 1)/2)^3 + \frac{n}{\kappa} (n + (n - 1)/2)^2 \right)\right), \quad (44)$$

where the first term again corresponds to the cost of LU decomposition, and the second term to the cost of

back substitutions. We can now bound N_s in terms of N_0 and κ . To do so, we need only realize that

$$N_s \approx N_0 \frac{\kappa}{n}, \quad (45)$$

where κ is typically closer to 1 than to n , making the constant $\frac{\kappa}{n}$ quite small. Thus, the assembly cost can be rewritten as

$$T_{assembly} = O\left(\frac{\kappa}{n} N_0 \left((n + (n-1)/2)^3 + \frac{n}{\kappa} (n + (n-1)/2)^2 \right)\right), \quad (46)$$

which is significantly smaller than if $\kappa = n$ as in the standard RBF-FD method. In practice, as n and ℓ increase, κ decreases because larger stencils allow us to retain more weights per stencil [49]. The only difference between steps 5 and 6 is that the values of n and ℓ are potentially different for the Laplacian versus point evaluation.

Steps 5 and 6 use Algorithm 1, which also require k-d trees to be built on the node sets for a cost of $O(N_0 \log N_0)$. The cost of searching them for n nearest neighbors on N_s stencils is $O(n N_s \log N_0)$. Using (45), we can write the total preprocessing cost as

$$T_{preprocessing} = O\left(N_\Gamma N_0 + \left((n + (n-1)/2)^3 + \frac{n}{\kappa} (n + (n-1)/2)^2 \right) \frac{\kappa}{n} N_0 + N_0 \log N_0 + \kappa N_0 \log N_0\right), \quad (47)$$

where again κ is a number closer to 1 than to n . We have observed in practice that the second term dominates this cost and scales as $O(N_0)$ for a given n and κ . To control efficiency, κ could be explicitly introduced as an input to Algorithm 1, but we leave this approach for future work.

5.2. Time-stepping complexity

We now estimate the complexity for steps 8-20 which are carried out in each timestep. We ignore steps 8, 9, 18, 19, and 20, as they are trivially estimated. Instead, we focus on the cost of a single time-step. Step 10 is done using the RK3 method, which has three stages, but the cost only depends on the number of seed nodes and the total number of embedded boundaries. This cost is clearly $O(N_\Gamma N_d)$. Step 11 is the same as the preprocessing steps 3 and 4, and therefore has a cost of $O(N_\Gamma N_0)$.

Steps 12 and 13 involve a mix of *copying* old information and *computing* new information, as shown in Algorithm 3. Estimating the cost of these steps requires an analysis of Algorithm 3, combined with a slightly modified analysis for Algorithm 1. Within Algorithm 3, letting $N_e = |X_e(t)|$, a k-d tree is first formed on $X_e(t)$ for $O(N_e \log N_e)$ operations in step 1. Steps 2 and 13 can be ignored as this can simply be done with Boolean flags. Steps 6, 10, and 12 can be done in constant time. This leaves the following steps:

- Step 4 can be done in $O(\log N_e)$ operations (with possible early termination).
- Step 5 has a cost of $O(n \log N_e)$.
- Step 7 can be done in practice by computing $\|p_{new} - p_{old}\|_2$, where p_{new} are the positions of neighbors at t_{n+1} , and p_{old} the positions at t_n ; this costs $O(n)$ operations. Steps 8 and 9 can be ignored.
- Step 11 involves a kd-tree search for a cost of $O(n \log N_e)$.

Each of these steps could, in the worst case, be executed N_s times (once for each stencil center). At the end of step 14, some fraction of the rows of $L(t_{n+1})$, $B(t_{n+1})$, and some of the local operators that constitute \mathcal{J}^{n+1} will have been computed. Let this fraction be τ , so that step 15 now only operates on $(1 - \tau)N_e$ nodes. The cost of step 15 can therefore be written based on the analysis from Section 5.1 as

$$T_{step\ 15} = O\left(\left((n + (n-1)/2)^3 + \frac{n}{\kappa} (n + (n-1)/2)^2 \right) \frac{\kappa}{n} (1 - \tau)N_e + \kappa(1 - \tau)N_e \log N_e\right). \quad (48)$$

The total cost of Algorithm 3 can thus be written using the above list, (45), and (48) as

$$T_{updates} = O\left(\left(1 + \frac{\kappa}{n} + 2\kappa\right) N_e \log N_e + \left(3\frac{\kappa}{n} + \kappa\right) N_e\right) + T_{step\ 15}. \quad (49)$$

In practice, $\tau \approx 1$ and $\kappa = O(1)$, allowing us to simply write

$$T_{updates} = O(N_e \log N_e), \quad (50)$$

i.e., k-d tree lookups dominate the cost of updating the differentiation matrices $L(t)$ and $B(t)$, and the interpolation operators $\mathcal{I}(t)$. In contrast, if τ is small, most of the rows of L and B , and most of the local interpolants comprising \mathcal{I} have to be recomputed. This concludes the analysis of steps 13 and 14 of Algorithm 2.

Steps 14 and 15 are the SL steps in the algorithm. Step 14 of Algorithm 2 involves 3 back traces costing $O(N_e)$, while step 15 implicitly involves a k-d tree lookup for each of the N_e points to find the correct local interpolation stencil, one back-solve on that stencil, and at least one evaluation on that stencil. In the worst case, we have $N_s = \frac{\kappa}{n} N_e$ stencils, exactly one back substitution per stencil for a cost of $O((n + (n - 1)/2)^2)$, and exactly one evaluation per stencil also for a cost of $O((n + (n - 1)/2)^2)$. The lookup costs $O(N_e \log N_e)$ operations. The total cost of Steps 14 and 15 of Algorithm 2 is therefore

$$T_{SL} = O\left(N_e \log N_e + 2\frac{\kappa}{n} N_e (n + (n - 1)/2)^2\right), \quad (51)$$

which is dominated by the second term. In practice, however, it is possible for our batching technique to result in multiple evaluations per stencil and fewer back-substitutions.

Step 16 of Algorithm 2 involves forming the preconditioner as outlined in Section 4.1. The cost of this step is dominated by the cost of forming the approximate Schur complement, which in turn involves sparse matrix multiplications of matrices with at most n non-zero entries per row. This incurs a cost of $O(n^2 N_e)$. Finally, the complexity of step 17 of Algorithm 2 is non-deterministic as it involves the convergence of the preconditioned GMRES method bootstrapped with a guess. Rather than attempt to estimate the complexity, we show results in terms of number of GMRES iterations in Section 6. Thus, the total cost of each time-step of Algorithm 2 (assuming $\tau \approx 1$) is $T_{step} = T_{updates} + T_{SL} + O(n^2 N_e) + \text{Cost of GMRES}$.

$$\implies T_{step} = O\left(2N_e \log N_e + 2\frac{\kappa}{n} N_e (n + (n - 1)/2)^2 + n^2 N_e\right) + \text{Cost of GMRES}. \quad (52)$$

Depending on which terms dominate above, this cost is either linear or log-linear in N_e (provided the cost of GMRES is kept low with a good preconditioner and initial guess). In Section 6.3, we demonstrate that the complexity of our method is indeed very close to linear in both the preprocessing and per-step costs, with a very large improvement in per-step costs (over the preprocessing costs) due to the reuse of weights over several steps.

6. Numerical Results

We test our numerical framework via convergence studies on the forced advection-diffusion equation with two Peclet numbers: 1 and 1000. The forcing term is selected to maintain a prescribed solution for all time, and the prescribed solution is used to test spatial convergence rates. We solve this test problem on irregular 2D and 3D domains with boundaries moving at the fluid velocity. Given a true solution $c(\mathbf{x}, t)$ and a numerical solution $C(\mathbf{x}, t)$, we compute relative ℓ_2 errors at the final time $t = 0.5$ on the node set X as $e_{\ell_2} = \frac{\|c_X - C_X\|_2}{\|c_X\|_2}$. In addition to relative ℓ_2 errors, we also report the average number of GMRES iterations per step for each value of the node count N and approximation order ξ . Finally, we verify the complexity estimates in Section 6 via timings, and also show a comparison of computational cost and accuracy as a function of the approximation order. In all cases, we set the GMRES tolerance to $\min(0.1h^\xi, 10^{-7})$.

6.1. Forced advection-diffusion in a time-varying 2D domain

The first test involves solving the advection-diffusion equation on an irregular time-varying 2D domain. The initial domain is the unit disk with two embedded ellipses defined in parametric form as:

$$E_1 : x = 0.4 \cos(\mu), y = -0.5 + 0.2 \sin(\mu), \quad (53)$$

$$E_2 : x = 0.1 \cos(\mu), y = 0.2 \sin(\mu), \quad (54)$$

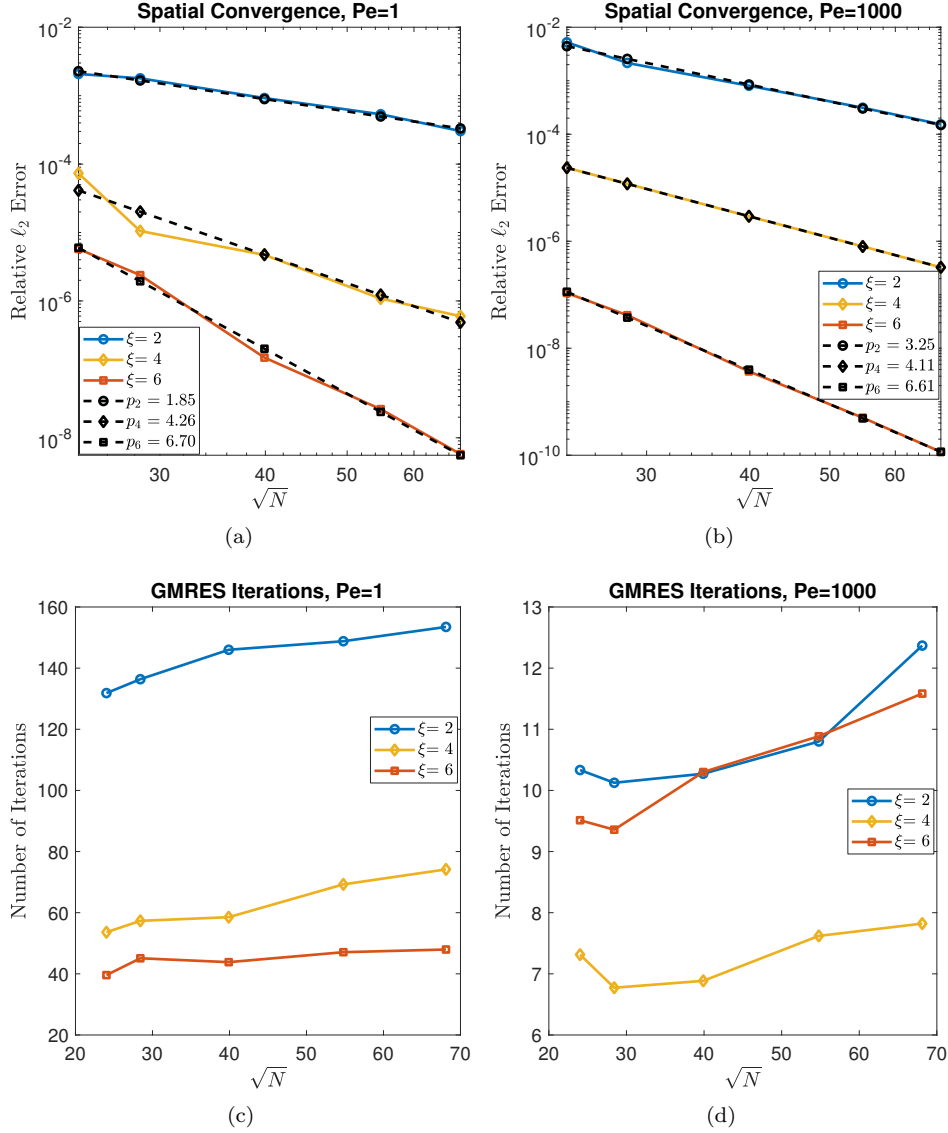


Figure 3: Top row: relative ℓ_2 error vs \sqrt{N} as a function of approximation order ξ for forced advection-diffusion on the unit disk. The dashed lines are lines of best fit indicating the slope (and hence convergence rate). Bottom row: average number of GMRES iterations per time-step as a function of \sqrt{N} and ξ .

where $-\pi \leq \mu < \pi$. The initial simulation domain is then given as $\Omega(t_0) = \{\mathbb{B}^2 \setminus (E_1 \cup E_2)\}$, *i.e.*, the portion of the domain outside the ellipses but within the unit disk. We use 20 seed nodes on the boundary of each ellipse, and reconstruct the moving boundary from these seed nodes using the (parametric) periodic degree-7 polyharmonic spline RBF within the geometric model developed by the authors [51]. The manufactured solution we use is given by

$$c(\mathbf{x}, t) = 1 + \sin(\pi x) \cos(\pi y) \sin(\pi t), \quad (55)$$

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

$$\mathbf{u}(\mathbf{x}, t) = \sin(\pi \|\mathbf{x}\|_2^2) \sin(\pi t) [y, -x]. \quad (56)$$

This velocity field vanishes on the boundary of the disk, and the 20 seed nodes on the embedded domain boundaries are advected with the velocity \mathbf{u} , thereby leading to deformation of the embedded ellipses. We

use a pure Neumann boundary condition operator \mathcal{B} ($\alpha = -\nu$ and $\beta = 0$), and the right hand side $g(\mathbf{x}, t)$ is given by applying \mathcal{B} to the prescribed $c(\mathbf{x}, t)$.

We measure errors in our numerical solution against the prescribed c . To obtain Peclet numbers $Pe = 1$ and $Pe = 1000$, we set $\nu = 1$ and $\nu = 10^{-3}$, respectively. We set the time-step as mentioned previously. We simulate the PDE to time $t = 0.5$ using Algorithm 2.

The results for $\xi = 2, 4, 6$ are shown in Figures 3(a) and 3(b), plotted as a function of \sqrt{N} (proportional to $1/h$). The results show that our predicted spatial convergence rate of h^ξ roughly holds under refinement for both Peclet numbers $Pe = 1$ and $Pe = 1000$; however, the errors are much lower for $Pe = 1000$. In addition, Figures 3(c) and 3(d) show the average number of GMRES iterations per timestep for each of these simulations, again as a function of \sqrt{N} . For $Pe = 1$, increasing ξ results in fewer GMRES iterations, possibly due to the interplay of ξ with the guess and preconditioner. For $Pe = 1000$, the iteration counts are more erratic as ξ is increased, but it is clear that far fewer GMRES iterations are required for this value of Pe , since a smaller value of ν was used (improving the conditioning of the time-stepping matrix).

6.2. Forced advection-diffusion in a time-varying 3D domain

Next, we do a convergence study on the forced advection-diffusion equation in an irregular time-varying 3D domain. The initial domain is the **unit ball** with an embedded ball E_3 of radius 0.2 centered at $(0.1, 0.2, 0.3)$. The initial simulation domain is $\Omega(t_0) = \{\mathbb{B}^3 \setminus E_3\}$. We use 200 seed nodes on the boundary of E_3 , and reconstruct the moving surface from these seed nodes using the (parametric) spherical, degree-8, polyharmonic spline RBF [51]. In this case, our manufactured solution is

$$c(\mathbf{x}, t) = 1 + \sin(\pi x) \cos(\pi y) \cos(\pi z) \sin(\pi t), \quad (57)$$

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

$$\mathbf{u}(\mathbf{x}, t) = \sin(\pi \|\mathbf{x}\|_2^2) \sin(\pi t) [yz, -2xz, xy]. \quad (58)$$

The boundary conditions, time-steps, and Peclet numbers are chosen as in the 2D case, and the seed nodes are once again moved with the local fluid velocity, causing deformation while enforcing a no-slip condition. We simulate the PDE to time $t = 0.5$, and measure errors against the manufactured solution.

The results are shown in Figure 4, plotted as function of $\sqrt[3]{N}$ (proportional to $1/h$). Once again, from Figures 4(a) and 4(b), we see that our results match the predicted spatial convergence rate of h^ξ except for the slightly erratic convergence at $Pe = 1000$ for $\xi = 2$ (likely due to an insufficiently small timestep on the finer node set). Figures 4(c) and 4(d) show the average number of GMRES iterations per timestep for each of those Peclet numbers as a function of $\sqrt[3]{N}$. In both cases, increasing ξ decreases the number of iterations (albeit erratically at low Peclet number). In addition, the number of iterations increase more slowly with $\sqrt[3]{N}$ as ξ is increased, demonstrating the efficiency of higher order methods. As in the 2D case, Figure 4(d) shows that a higher Peclet number requires fewer GMRES iterations to obtain the same tolerance, once again because of the improved conditioning due to smaller values of the diffusion coefficient ν .

6.3. Timings

Next, we present timing results to verify the complexity estimates of our method. We present timings only for $Pe = 1000$. This is because our goal is to verify the analysis from Section 5, which only abstractly includes the cost of the GMRES method. Since the results thus far have already shown that the number of GMRES iterations is smaller for $Pe = 1000$, this case will be more ideal for verifying complexity. We present two types of results. All timings were run on a PC with an 8-core Intel i7-9700K CPU clocked at a base speed of 3.6 GHz which had 32 GB of 2.6GHz RAM. Timings were accomplished in Matlab R2020b without any explicit parallelization of the code.

In Figure 5, we present wall-clock times as a function of the number of nodes N and the approximation order ξ for both the 2D and 3D test problems. We separate the timings into preprocessing time (Figures 5(a) and (c)) and average per-step time (Figures 5(b) and (d)). The lines of best fit indicate that the slopes are close to 1, which indicates linear or quasi-linear complexity. In addition, it is clear that the cost per time-step is

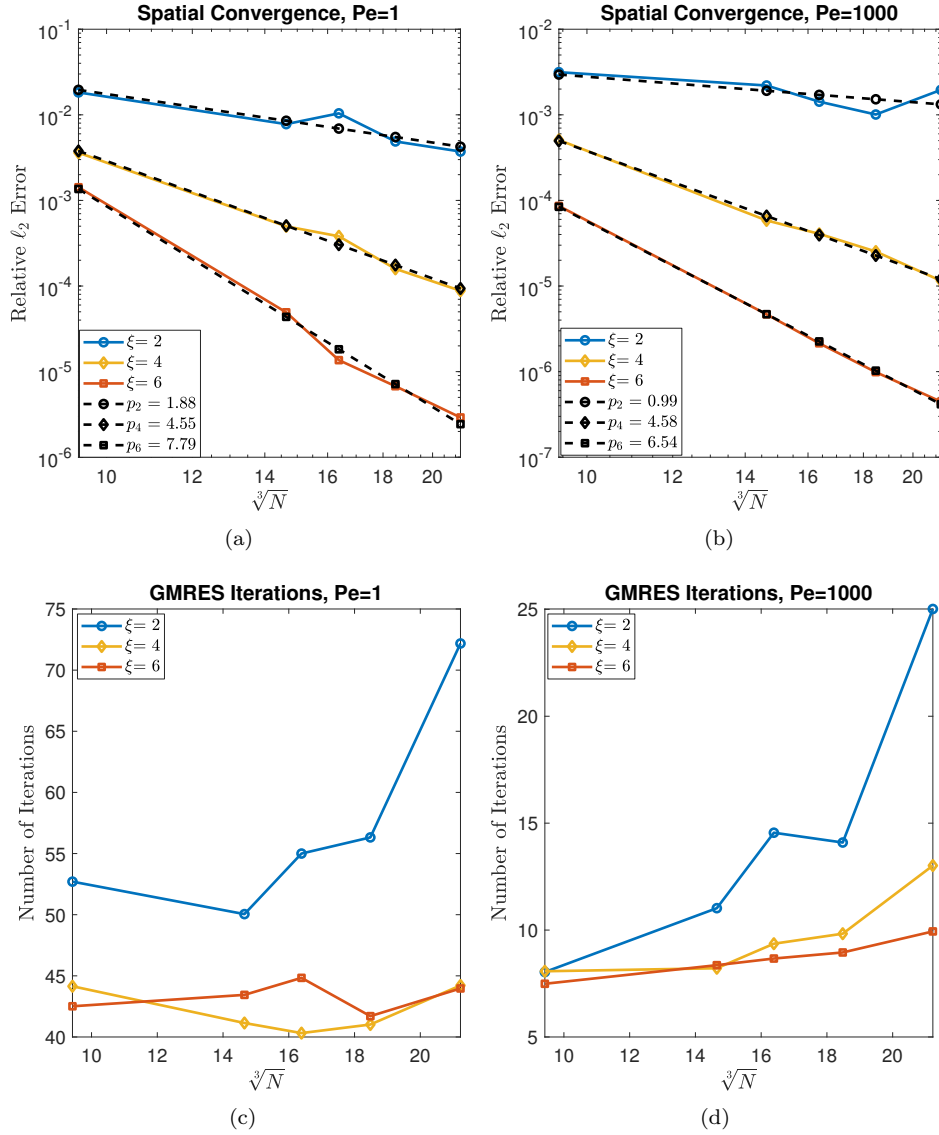


Figure 4: Top row: relative ℓ_2 error vs $\sqrt[3]{N}$ as a function of approximation order ξ for forced advection-diffusion in the unit ball. The dashed lines are lines of best fit indicating the slope (and hence convergence rate). Bottom row: average number of GMRES iterations per time-step as a function of $\sqrt[3]{N}$ and ξ .

much lower than the preprocessing cost due to our strategy of copying as many RBF-FD weights as possible. The difference in costs is seen particularly in Figures 5(c) and (d) (for the 3D simulation), which show a full order of magnitude difference in time between a single preprocessing step and the average time-step. Even more interestingly, for the 3D results, we see that the sixth-order method ($\xi = 6$) is less expensive than the fourth-order method ($\xi = 4$) not just in terms of efficiency but in terms of wall-clock time. This is likely because the sixth-order method allows for a greater number of stable weights to be computed and retained per stencil, consequently allowing fewer stencils and recomputations in total.

In Figure 6, we present the cost-accuracy tradeoffs for the 2D and 3D tests, again as a function of number of nodes N and ξ . Here, we focus purely on the per-step costs, since the preprocessing costs are not important for long-running simulations. The figures all show that the higher order methods deliver greater accuracy for the same wall-clock time (and a given value of N), since our manufactured solutions are smooth. In addition, we once again see in Figure 6b that the sixth-order method takes even less time than the fourth-

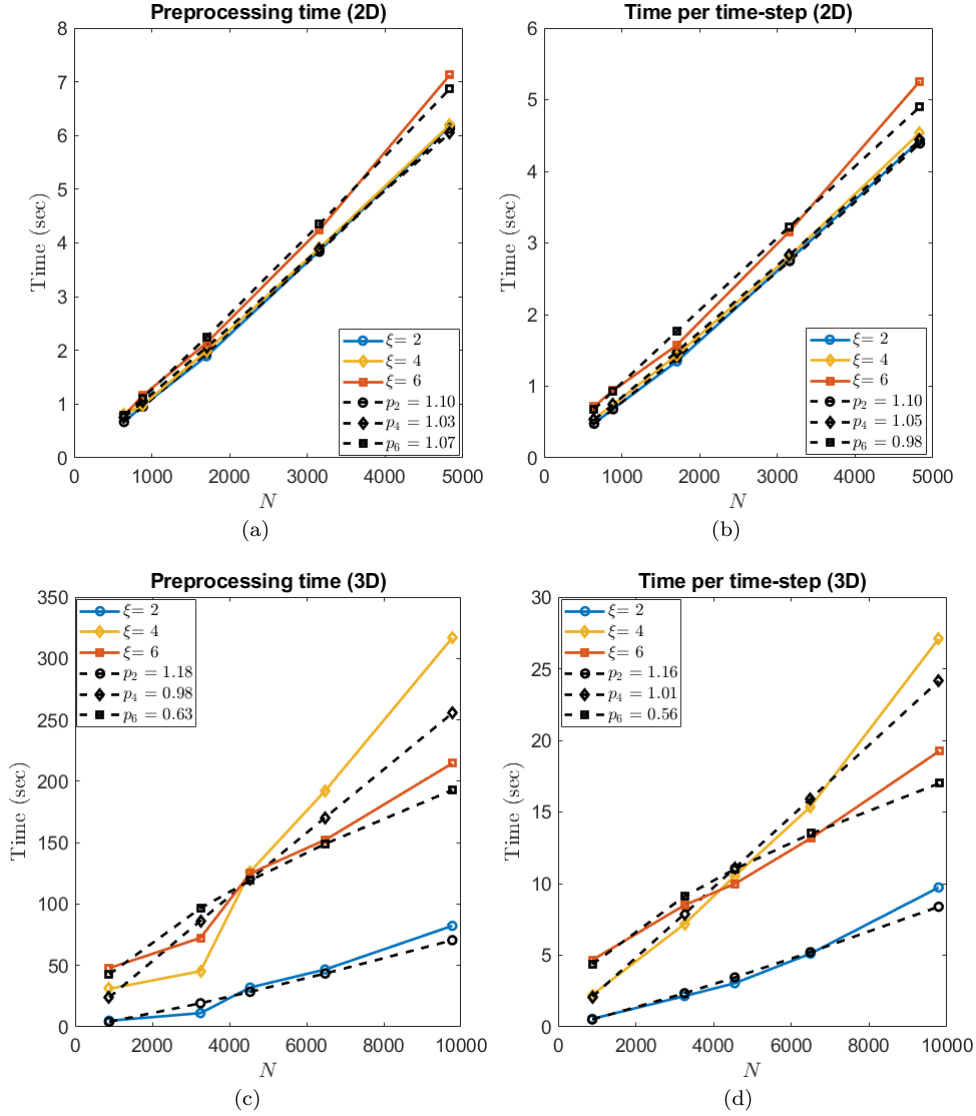


Figure 5: Top row: wall-clock time vs N as a function of approximation order ξ for forced advection-diffusion on the unit disk. The dashed lines are lines of best fit indicating the slope (and hence computational complexity). Bottom row: wall-clock time vs N as a function of order ξ for forced advection-diffusion in the unit ball. All timings are for Peclet number $Pe = 1000$.

order method for the largest value of N (top-most points of the curves) while delivering greater accuracy (in 3D). Both these timing tests confirm the quasilinear complexity of our method while also demonstrating the benefits of using high-order methods especially in higher dimensions. These tests also clearly demonstrate the importance of selective updates to RBF-FD weights and stencils using Algorithm 3.

7. A coupled problem

To fully demonstrate the utility of our numerical methods, we now apply them to solving a 3D coupled problem. In this problem, we track a chemical concentration $c(\mathbf{x}, t)$ in a fluid inside the same initial domain as in Section 6.2, with the embedded boundary passively advected by an incompressible velocity field $\mathbf{u}(\mathbf{x}, t)$. The concentration $c(\mathbf{x}, t)$ therefore satisfies the advection-diffusion equation (written using the




Figure 6: Cost vs accuracy as a function of approximation order ξ for the 2D problem (left) and the 3D problem (right). The figure shows average time per time-step and the relative ℓ_2 error as N is increased for a given value of ξ . All timings are for Peclet number $Pe = 1000$.

Lagrangian/material derivative):

$$\frac{dc}{dt} = \nu \Delta c + f_1(\mathbf{x}, t), \mathbf{x} \in \Omega(t), \quad (59)$$

where $f_1(\mathbf{x}, t)$ is some forcing term, and $\Omega(t)$ is the time-varying irregular domain defined by the outer spherical boundary and the deforming inner embedded boundary $\Gamma(t)$. On the boundary of the unit ball, c satisfies a time-varying inhomogeneous Neumann boundary condition:

$$-\nu \frac{\partial c(\mathbf{x}, t)}{\partial \mathbf{n}} = g(\mathbf{x}, t), \mathbf{x} \in \mathbb{S}, \quad (60)$$

which corresponds to an *inward* flux of c . The inner moving embedded boundary $\Gamma(t)$ is viewed as a reactive “zone” on which chemicals can bind, unbind, and participate in other reactions. We track the *bound* chemical surface density separately, and label it C_B . We assume that at each point on the embedded boundary $\Gamma(t)$, C_B satisfies the following PDE:

$$\frac{\partial C_B}{\partial t} + \mathbf{u} \cdot \nabla C_B - C_B \nabla_{\Gamma} \cdot \mathbf{u} = k_{on}(C^{Tot} - C_B) - k_{off}C_B + k_{self}C_B(C^{Tot} - C_B) + f_2, \quad (61)$$

where k_{on} and k_{off} are the binding and unbinding rates of C_B , k_{self} is the rate at which C_B reacts with itself, C^{Tot} is the total density of binding sites at each point of the reactive zone, f_2 is some forcing term, and $\nabla_{\Gamma} \cdot$ is the surface divergence operator. The last term on the left accounts for local surface changes due to the flow. As before, the first two terms in (61) can be combined into the material derivative, yielding the following equation for C_B :

$$\frac{dC_B}{dt} - C_B \nabla_{\Gamma} \cdot \mathbf{u} = k_{on}(C^{Tot} - C_B) - k_{off}C_B + k_{self}C_B(C^{Tot} - C_B) + f_2, \quad (62)$$

Balancing fluxes at the interface $\Gamma(t)$ yields the following time-varying Robin boundary condition on $c(\mathbf{x}, t)$:

$$-\nu \frac{\partial c(\mathbf{x}, t)}{\partial \mathbf{n}} = -k_{on}(C^{Tot} - C_B) + k_{off}C_B, \mathbf{x} \in \Gamma(t). \quad (63)$$

This model represents a *one-way* coupled bulk-surface problem, and is loosely inspired by similar problems arising in the context of platelet aggregation and coagulation.

7.1. A manufactured solution to the coupled problem

We specify the concentration $c(\mathbf{x}, t)$ to be

$$c(\mathbf{x}, t) = c(x, y, z, t) = 1 + \sin(\pi x) \cos(\pi y) \sin(\pi z) \sin(\pi t). \quad (64)$$

In addition, we *prescribe* an incompressible velocity field $\mathbf{u}(\mathbf{x}, t) = (u, v, w)$ in the unit ball as

$$\mathbf{u}(\mathbf{x}, t) = \sin(\pi t) \sin(\pi \|\mathbf{x}\|_2^2) [yz, -2xz, xy]. \quad (65)$$

Because points on the embedded boundary $\Gamma(t)$ will advect in this velocity field, the no-slip boundary is automatically satisfied on its surface. Then, we set the forcing term $f_1(\mathbf{x}, t)$ as

$$f_1(\mathbf{x}, t) = \frac{\partial c}{\partial t} + \mathbf{u} \cdot \nabla c - \nu \Delta c. \quad (66)$$

The boundary condition function $g(\mathbf{x}, t)$ is then obtained by applying the Neumann operator $-\nu \frac{\partial}{\partial \mathbf{n}}$ to $c(\mathbf{x}, t)$. We substitute $c(\mathbf{x}, t)$ into (63) and solve for C_B to obtain

$$C_B = \frac{-\nu \frac{\partial c}{\partial \mathbf{n}} + k_{on} C^{Tot}}{k_{on} + k_{off}}. \quad (67)$$

Using C_B , we compute the forcing term f_2 as

$$f_2 = \frac{dC_B}{dt} - C_B \nabla_{\Gamma} \cdot \mathbf{u} - k_{on} (C^{Tot} - C_B) c_{amb} + k_{off} C_B - k_{self} C_B (C^{Tot} - C_B). \quad (68)$$

For a spatial convergence study, we define numerical errors by comparing against c and C_B . Since \mathbf{u} is known analytically, the term ∇_{Γ} can be computed quasi-analytically using the numerically-computed normal vectors on $\Gamma(t)$.

7.2. Time-stepping for the coupled problem

The fluid-phase chemicals $c(\mathbf{x}, t)$ and the bound chemical density C_B are coupled through the boundary conditions in (63), which in turn reflect the binding and unbinding reactions in the ODE (61). To simulate this system efficiently, we use a simple time-splitting scheme that is a combination of SL updates for c and full Lagrangian updates for C_B . Given C_B^n , c^n , and locations for the seed nodes on Γ^n , the algorithm is as follows:

1. Solve (62) as an ODE using the semi-implicit BDF3 (SBDF3) method [1] to advance C_B^n to \tilde{C}_B . The forcing term f_2 is treated implicitly in time, and all other terms are treated explicitly.
2. Advance the locations of the seed nodes on Γ^n using $\mathbf{u}(\mathbf{x}, t)$ and the RK3 method. This advects \tilde{C}_B in a Lagrangian fashion to obtain C_B^{n+1} . Use the new seed node locations to construct a geometric representation of the embedded boundary and to interpolate C_B to boundary nodes.
3. Use this representation and the interpolated C_B^{n+1} to obtain the boundary conditions for c^{n+1} on Γ^{n+1} as:

$$-\nu \frac{\partial c^{n+1}}{\partial \mathbf{n}} = -k_{on} (C^{Tot} - C_B^{n+1}) + k_{off} C_B^{n+1}. \quad (69)$$

4. Update c^n to c^{n+1} using the SL method presented in Algorithm 2.

It is important to note that a fully-coupled problem in which c participates in the first term on the right hand side of (62) would likely require high-order temporal splitting; we leave such an investigation to future work. For the purposes of this article, we have observed that the above scheme produces highly accurate results both for c and C_B . It is also worth noting that if \mathbf{u} is not known analytically, the $\nabla_{\Gamma} \cdot \mathbf{u}$ term must be computed numerically. This can also be done using a stabilized version of overlapped RBF-FD specialized to manifolds [52, 57].

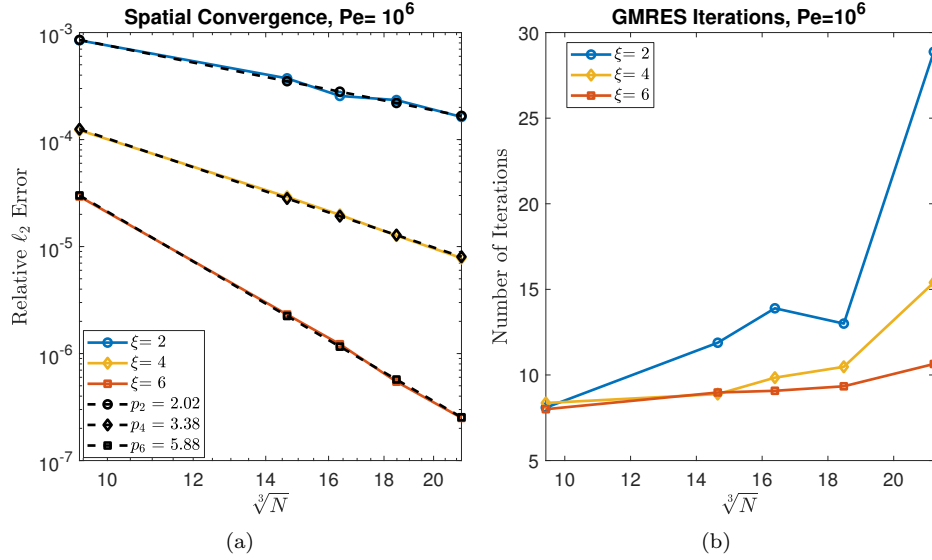


Figure 7: Left: relative errors vs node spacing as a function of approximation order ξ . Right: number of GMRES iterations as a function of ξ .

7.3. Results

To more fully test the stability of our method, we now set the diffusion coefficient to $\nu = 10^{-6}$, and scale the velocity field \mathbf{u} to obtain a Peclet number of $Pe = 10^6$. We select the time-step as in the advection-diffusion test cases, run convergence studies to time $t = 0.5$ using $\xi = 2, 4$ and 6 , and measure both the relative error in c and the average number of GMRES iterations. The results are shown in Figure 7. Much like our other tests, Figure 7(a) shows that the spatial error in the numerical approximation to c decreases at the rate of approximately h^ξ . In addition, Figure 7(b) shows that the number of GMRES iterations increases approximately linearly with the node spacing, and decreases as ξ is increased. It is also worth noting that the relative error in C_B is purely temporal and is on the order of 10^{-8} for all the values of Δt used in this test. This application clearly demonstrates the utility of our overall framework in solving coupled bulk-surface problems on moving domains.

8. Summary and Future Work

In this article, we presented a high-order numerical method for simulating the advection-diffusion equation on domains with time-varying embedded boundaries. Our method uses semi-Lagrangian (SL) advection combined with an Eulerian formulation for diffusion, both applied in the context of a rapid node adaptation algorithm. These techniques rely on a generalization of the overlapped RBF-FD method that replaces the hand-tuned overlap parameter δ with a pair of automatically-computed stability indicators. We also presented a novel automatic algorithm for updating RBF-FD interpolation stencils and differentiation matrices on a time-varying node set. We conducted an informal error analysis to show the high-order convergence rate of our method, and conducted a computational complexity analysis to determine the efficiency of our algorithms. We verified the aforementioned high-order convergence rates on both 2D and 3D advection-diffusion problems on irregular domains with moving embedded boundaries. In addition, we demonstrated high-order convergence rates on a more complicated 3D coupled problem.

While our informal error analysis provided estimates that were verified numerically, a formal error analysis of semi-Lagrangian RBF-FD methods for advection-diffusion equations is absent in the literature. We plan to address this issue in future work. In addition, in order to apply our methods to problems requiring very large node sets, the overlapped RBF-FD method (and the overall method presented in this article) must be parallelized for distributed and shared memory architectures. We plan to tackle this in future work as well.

Acknowledgments

VS was supported by NSF grants CISE CCF 1714844 and DMS-1521748. ALF was supported by NSF grant DMS-1521748 and NIHBL grant 1U01HL143336. GBW was supported by NSF grants CISE CCF 1717556 and DMS 1952674.

References

References

- [1] Ascher, U. M., Ruuth, S. J., and Wetton, B. T. R. (1997). Implicit-Explicit Methods For Time-Dependent PDEs. *SIAM J. Numer. Anal.*, 32:797–823.
- [2] Barnett, G. A. (2015). *A Robust RBF-FD Formulation based on Polyharmonic Splines and Polynomials*. PhD thesis, University of Colorado Boulder.
- [3] Bayona, V. (2019). Comparison of moving least squares and RBF+ poly for interpolation and derivative approximation. *Journal of Scientific Computing*, 81(1):486–512.
- [4] Bayona, V., Flyer, N., and Fornberg, B. (2019). On the role of polynomials in rbf-fd approximations: Iii. behavior near domain boundaries. *Journal of Computational Physics*, 380:378–399.
- [5] Bayona, V., Flyer, N., Fornberg, B., and Barnett, G. A. (2017). On the role of polynomials in RBF-FD approximations: II. Numerical solution of elliptic PDEs. *J. Comput. Phys.*, 332:257–273.
- [6] Bayona, V., Moscoso, M., Carretero, M., and Kindelan, M. (2010). RBF-FD formulas and convergence properties. *J. Comput. Phys.*, 229(22):8281–8295.
- [7] Behrens, J. and Iske, A. (2002). Grid-free adaptive semi-Lagrangian advection using radial basis functions. *Comput. Math. Appl.*, 43(3):319–327.
- [8] Benzi, M., Golub, G. H., and Liesen, J. (2005). Numerical solution of saddle point problems. *Acta Numerica*, 14:1–137.
- [9] Benzi, M. and Wathen, A. J. (2008). Some preconditioning techniques for saddle point problems. In *Model order reduction: theory, research aspects and applications*, pages 195–211. Springer.
- [10] Bonaventura, L. and Ferretti, R. (2014). Semi-Lagrangian methods for parabolic problems in divergence form. *SIAM Journal on Scientific Computing*, 36(5):A2458–A2477.
- [11] Bonaventura, L. and Ferretti, R. (2016). Flux form semi-Lagrangian methods for parabolic problems. *Communications in Applied and Industrial Mathematics*, 7(3):56–73.
- [12] Bowman, J. C., Yassaie, M. A., and Basu, A. (2015). A fully Lagrangian advection scheme. *Journal of Scientific Computing*, 64(1):151–177.
- [13] Calhoun, D. and LeVeque, R. J. (2000). A Cartesian grid finite-volume method for the advection-diffusion equation in irregular geometries. *J. Comput. Phys.*, 157(1):143–180.
- [14] Colonius, T. and Taira, K. (2008). A fast immersed boundary method using a nullspace approach and multi-domain far-field boundary conditions. *Compu. Methods Appl. Mech. Engrg.*, 197:2131–2146.
- [15] Davydov, O. and Oanh, D. T. (2011). Adaptive meshless centres and RBF stencils for Poisson equation. *J. Comput. Phys.*, 230(2):287–304.
- [16] Davydov, O. and Schaback, R. (2018). Minimal numerical differentiation formulas. *Numerische Mathematik*, 140(3):555–592.
- [17] Degond, P. and Mas-Gallic, S. (1989). The weighted particle method for convection-diffusion equations. part 1: The case of an isotropic viscosity. *Mathematics of Computation*, 53(188):485–507.

- [18] Duff, I. S. and Koster, J. (2001). On algorithms for permuting large entries to the diagonal of a sparse matrix. *SIAM Journal on Matrix Analysis and Applications*, 22(4):973–996.
- [19] Fadlun, E., Verzicco, R., Orlandi, P., and Mohd-Yusof, J. (2000). Combined immersed-boundary finite-difference methods for three-dimensional complex flow simulations. *J. Comput. Phys.*, 161:35–60.
- [20] Falcone, M. and Ferretti, R. (1998). Convergence analysis for a class of high-order semi-Lagrangian advection schemes. *SIAM Journal on Numerical Analysis*, 35(3):909–940.
- [21] Falcone, M. and Ferretti, R. (2013). *Semi-Lagrangian approximation schemes for linear and Hamilton—Jacobi equations*. SIAM.
- [22] Fasshauer, G. E. (2007). *Meshfree Approximation Methods with MATLAB*. Interdisciplinary Mathematical Sciences - Vol. 6. World Scientific Publishers, Singapore.
- [23] Flyer, N., Barnett, G. A., and Wicker, L. J. (2016a). Enhancing finite differences with radial basis functions: Experiments on the Navier-Stokes equations. *J. Comput. Phys.*, 316:39–62.
- [24] Flyer, N., Fornberg, B., Bayona, V., and Barnett, G. A. (2016b). On the role of polynomials in RBF-FD approximations: I. Interpolation and accuracy. *J. Comput. Phys.*, 321:21–38.
- [25] Flyer, N., Lehto, E., Blaise, S., Wright, G. B., and St-Cyr, A. (2012). A guide to RBF-generated finite differences for nonlinear transport: shallow water simulations on a sphere. *J. Comput. Phys.*, 231:4078–4095.
- [26] Flyer, N. and Wright, G. B. (2007). Transport schemes on a sphere using radial basis functions. *J. Comput. Phys.*, 226:1059–1084.
- [27] Flyer, N. and Wright, G. B. (2009). A radial basis function method for the shallow water equations on a sphere. *Proc. Roy. Soc. A*, 465:1949–1976.
- [28] Fornberg, B. and Lehto, E. (2011). Stabilization of RBF-generated finite difference methods for convective PDEs. *J. Comput. Phys.*, 230:2270–2285.
- [29] Fuselier, E. J. and Wright, G. B. (2013). A high-order kernel method for diffusion and reaction-diffusion equations on surfaces. *J. Sci. Comput.*, 56(3):535–565.
- [30] Glowinski, R., Pan, T., and Périaux, J. (1998). Distributed Lagrange multiplier methods for incompressible viscous flow around moving rigid bodies. *Comput. Methods Appl. Mech. Engrg.*, 151:181–194.
- [31] Goldstein, D., Handler, R., and Sirovich, L. (1993). Modeling a no-slip flow boundary with an external force field. *J. Comput. Phys.*, 105:354–366.
- [32] Gritton, C., Guilkey, J., Hooper, J., Bedrov, D., Kirby, R. M., and Berzins, M. (2017). Using the material point method to model chemical/mechanical coupling in the deformation of a silicon anode. *Modelling and Simulation in Materials Science and Engineering*, 25(4):045005.
- [33] Johansen, H. S. and Colella, P. (1998). A Cartesian grid embedded boundary method for poisson’s equation on irregular domains. *J. Comput. Phys.*, 147:60–85.
- [34] Kim, J., Kim, D., and Choi, H. (2001). An immersed-boundary finite-volume method for simulations of flow in complex geometries. *J. Comput. Phys.*, 171:132–150.
- [35] Le Roux, D. Y., Lin, C. A., and Staniforth, A. (1997). An accurate interpolating scheme for semi-Lagrangian advection on an unstructured mesh for ocean modelling. *Tellus A*, 49(1):119–138.
- [36] Lehto, E., Shankar, V., and Wright, G. B. (2017). A radial basis function (RBF) compact finite difference (FD) scheme for reaction-diffusion equations on surfaces. *SIAM J. Sci. Comput.*, 39:A2129–A2151.
- [37] Leiderman, K. and Fogelson, A. L. (2014). An Overview of Mathematical Modeling of Thrombus Formation Under Flow. *Thromb. Res.*, 133 Suppl:S12–S14.

- [38] Leiderman, K. M. and Fogelson, A. L. (2011). Grow with the flow: a spatial-temporal model of platelet deposition and blood coagulation under flow. *Math. Med. Biol.*, 28:47–84.
- [39] Leiderman, K. M. and Fogelson, A. L. (2013). The influence of hindered transport on the development of platelet thrombi under flow. *Bull. of Math. Biol.*, 75:1255–1283.
- [40] LeVeque, R. J. and Li, Z. (1994). The immersed interface method for elliptic equations with discontinuous coefficients and singular sources. *SIAM J. Numer. Anal.*, 31:1001–1025.
- [41] McCorquodale, P., Colella, P., and Johansen, H. (2001). A Cartesian grid embedded boundary method for the heat equation on irregular domains. *J. Comput. Phys.*, 173:620–635.
- [42] Mohd-Yusof, J. (1997). Combined immersed-boundary/B-spline methods for simulations of flow in complex geometries. *Annu. Res. Briefs, Cent. Turbul. Res.*, pages 317–328.
- [43] Peskin, C. S. (1972). Flow pattern around heart valves: a numerical method. *J. Comput. Phys.*, 10:252–271.
- [44] Peskin, C. S. (1977). Numerical analysis of blood flow in the heart. *J. Comput. Phys.*, 25:220–252.
- [45] Peskin, C. S. (2002). The immersed boundary method. *Acta Numerica*, 11:479–517.
- [46] Piret, C. (2012). The orthogonal gradients method: A radial basis functions method for solving partial differential equations on arbitrary surfaces. *J. Comput. Phys.*, 231(20):4662–4675.
- [47] Piret, C. and Dunn, J. (2016). Fast RBF OGr for solving pdes on arbitrary surfaces. *AIP Conference Proceedings*, 1776(1).
- [48] Saad, Y. (2003). *Iterative methods for sparse linear systems*. SIAM.
- [49] Shankar, V. (2017). The overlapped radial basis function-finite difference (RBF-FD) method: A generalization of RBF-FD. *J. Comput. Phys.*, 342:211–228.
- [50] Shankar, V. and Fogelson, A. L. (2018). Hyperviscosity-based stabilization for radial basis function-finite difference (rbf-fd) discretizations of advection–diffusion equations. *J. Comput. Phys.*, 372:616 – 639.
- [51] Shankar, V., Kirby, R., and Fogelson, A. (2018a). Robust node generation for mesh-free discretizations on irregular domains and surfaces. *SIAM Journal on Scientific Computing*, 40(4):A2584–A2608.
- [52] Shankar, V., Narayan, A., and Kirby, R. M. (2018b). Rbf-loi: Augmenting radial basis functions (rbfs) with least orthogonal interpolation (loi) for solving pdes on surfaces. *Journal of Computational Physics*, 373:722–735.
- [53] Shankar, V. and Wright, G. B. (2018). Mesh-free semi-lagrangian methods for transport on a sphere using radial basis functions. *J. Comput. Phys.*, 366(C):170–190.
- [54] Shankar, V., Wright, G. B., Fogelson, A. L., and Kirby, R. M. (2014a). A radial basis function (RBF) finite difference method for the simulation of reaction-diffusion equations on stationary platelets within the augmented forcing method. *Inter. J. Numer. Methods Fluids*, 75(1):1–22.
- [55] Shankar, V., Wright, G. B., Kirby, R. M., and Fogelson, A. L. (2014b). A radial basis function (RBF)-finite difference (FD) method for diffusion and reaction–diffusion equations on surfaces. *J. Sci. Comput.*, 63(3):745–768.
- [56] Shankar, V., Wright, G. B., Kirby, R. M., and Fogelson, A. L. (2015). Augmenting the immersed boundary method with radial basis functions (rbfs) for the modeling of platelets in hemodynamic flows. *International Journal for Numerical Methods in Fluids*, 79(10):536–557.

- [57] Shankar, V., Wright, G. B., and Narayan, A. (2020). A robust hyperviscosity formulation for stable RBF-FD discretizations of Advection-Diffusion-Reaction equations on manifolds. *SIAM Journal on Scientific Computing*, 42(4):A2371–A2401.
- [58] Smolarkiewicz, P. K. and Margolin, L. G. (1997). On forward-in-time differencing for fluids: an Eulerian/semi-Lagrangian non-hydrostatic model for stratified flows. *Atmos.-Ocean*, 35:127–152.
- [59] Smolarkiewicz, P. K. and Pudykiewicz, J. A. (1992). A class of semi-Lagrangian approximations for fluids. *J. Atmos. Sci.*, 49(22):2082–2096.
- [60] Staniforth, A. and Côté, J. (1991). Semi-Lagrangian integration schemes for atmospheric models—a review. *Mon. Wea. Rev.*, 119(9):2206–2223.
- [61] Staniforth, A. and Wood, N. (2008). Aspects of the dynamical core of a nonhydrostatic, deep-atmosphere, unified weather and climate-prediction model. *J. Comput. Phys.*, 227(7):3445–3464.
- [62] Stein, D. B., Guy, R. D., and Thomases, B. (2016). Immersed boundary smooth extension: a high-order method for solving PDE on arbitrary smooth domains using fourier spectral methods. *J. Comput. Phys.*, 304:252–274.
- [63] Stein, D. B., Guy, R. D., and Thomases, B. (2017). Immersed boundary smooth extension (IBSE): A high-order method for solving incompressible flows in arbitrary smooth domains. *J. Comput. Phys.*, 335:155–178.
- [64] Trask, N., Maxey, M., Kim, K., Perego, M., Parks, M. L., Yang, K., and Xu, J. (2015). A scalable consistent second-order SPH solver for unsteady low Reynolds number flows. *Computer Methods in Applied Mechanics and Engineering*, 289:155–178.
- [65] Udaykumar, H. S., Mittal, R., and Shyy, W. (1999). Computation of solid-liquid phase fronts in the sharp interface limit on fixed grids. *J. Comput. Phys.*, 153:535–574.
- [66] Wendland, H. (2005). *Scattered data approximation*, volume 17 of *Cambridge Monogr. Appl. Comput. Math.* Cambridge University Press, Cambridge.
- [67] Wright, G. B. and Fornberg, B. (2006). Scattered node compact finite difference-type formulas generated from radial basis functions. *J. Comput. Phys.*, 212(1):99–123.
- [68] Xiu, D. and Karniadakis, G. E. (2001). A semi-Lagrangian high-order method for Navier–Stokes equations. *J. Comput. Phys.*, 172(2):658–684.
- [69] Yao, L. and Fogelson, A. L. (2012). Simulations of chemical transport and reaction in a suspension of cells I: an augmented forcing point method for the stationary case. *Inter. J. Numer. Methods Fluids*, 69(11):1736–1752.
- [70] Ye, T., Mittal, R., Udaykumar, H. S., and Shyy, W. (1999). An accurate Cartesian grid method for viscous incompressible flows with complex immersed boundaries. *J. Comput. Phys.*, 156:209–240.