Augmenting the immersed boundary method with Radial Basis Functions (RBFs) for the modeling of platelets in hemodynamic flows

Varun Shankar[1,*,†], Grady B. Wright[2], Robert M. Kirby[3] and Aaron L. Fogelson[4]

[1]*Department of Mathematics, University of Utah, Salt Lake City, UT, USA*
[2]*Department of Mathematics, Boise State University, Boise, ID, USA*
[3]*School of Computing, University of Utah, Salt Lake City, UT, USA*
[4]*Departments of Mathematics and Bioengineering, University of Utah, Salt Lake City, UT, USA*

SUMMARY

We present a new computational method by extending the immersed boundary (IB) method with a geometric model based on parametric radial basis function (RBF) interpolation of the Lagrangian structures. Our specific motivation is the modeling of platelets in hemodynamic flows, although we anticipate that our method will be useful in other applications involving surface elasticity. The efficacy of our new RBF-IB method is shown through a series of numerical experiments. Specifically, we test the convergence of our method and compare our method with the traditional IB method in terms of computational cost, maximum stable time-step size, and volume loss. We conclude that the RBF-IB method has advantages over the traditional IB method and is well-suited for modeling of platelets in hemodynamic flows. Copyright © 2015 John Wiley & Sons, Ltd.

1. INTRODUCTION

The immersed boundary (IB) Method was introduced by Charles Peskin in the early 1970s to solve the coupled equations of motion of a viscous, incompressible fluid and one or more massless, elastic surfaces or objects immersed in the fluid [1]. The IB method was originally developed to model blood flow in the heart and through heart valves [1–3], but has since been used in a wide variety of other applications, particularly in biofluid dynamics problems where complex geometries and immersed elastic membranes or structures are present and make traditional computational approaches difficult. Examples include swimming of organisms [4, 5], biofilm processes [6], mechanical properties of cells [7], cochlear dynamics [8], and insect flight [9, 10]. In this work, we are motivated by the application of the IB method to platelet aggregation in blood clotting, but expect our method to be useful in other applications as well.

Intravascular blood clots (thrombi) are initiated by damage to the endothelial cell lining of a blood vessel and involve the formation on the damaged surface of clumps of cells intermixed with a fibrous protein gel. The cells involved in this process are platelets, and the subject of this paper is a new approach to modeling platelets in order to simulate their adhesion to the injured vascular wall and cohesion with one another during the formation of a thrombus. The IB method is used to describe the mechanical interactions among a collection of discrete platelets, the background fluid, and the vessel wall. However, we now model our platelets with radial basis functions (RBFs) in order to achieve more accurate and less costly simulations.

---

*Correspondence to: Varun Shankar, Department of Mathematics, University of Utah, Salt Lake City, UT, USA.
†E-mail: vshankar@math.utah.edu

In this introduction, we briefly outline the relevant biology, describe how the IB method has been used in our previous platelet aggregation simulations, and give an overview of how use of the new method changes this description.

1.1. Modeling the mechanics of platelet aggregation

Disruption of the endothelial cell lining exposes collagen and adsorbed von Willebrand factor molecules in the subendothelial matrix to the blood. Platelets adhere to both molecules via specific receptor molecules on the platelets' surfaces. In addition to slowing or stopping platelet motion over the subendothelium, this binding triggers intracellular signaling pathways that lead to platelet activation [11, 12].

Our platelet aggregation models [4, 13–16] track the motion and behavior of a collection of individual platelets as they interact with the suspending fluid, one another, and the vessel walls. We also track fluid concentrations of platelet-activating chemicals, cell–cell and cell–surface forces, fluid motion, and the local fluid forces on the growing thrombus. In our models, nonactivated platelets are activated by proximity to reactive sites on the injured wall or through exposure to a sufficiently high concentration of activator in the fluid. Activation enables a platelet to cohere with other activated platelets and to secrete additional activator. The platelets and the secreted chemical move by advection with the fluid and diffusion relative to it. Each platelet is represented as an IB object, that is, as a collection of elastically linked Lagrangian points that each move at the local fluid velocity. New elastic links are created dynamically to model the adhesion of a platelet to the injured wall or the cohesion of activated platelets to one another. Multiple links can form between a pair of activated model platelets or between a model platelet and the injured wall, and these links collectively represent the ensemble of molecular bridges binding real platelets to one another or to the damaged vessel. The links exert forces on the surrounding fluid to resist motions that would otherwise separate the linked entities. Links may break if subject to sufficiently high stress. Model variables are fully coupled: the fluid carries the activator and platelets, while the interplatelet forces, potentiated by chemically induced activation of the platelets, determine the local flow. In this paper, we focus on mechanical interactions, not the activation process, and so we specify the conditions under which a platelet becomes activated and able to cohere with other activated platelets.

1.2. Motivation for the radial basis function–immersed boundary method

We model platelets as closed curves of interconnected IB points in 2D. A platelet's area or volume is determined by the region enclosed by the curve or surface and is preserved because of the incompressibility of the fluid. Inactive platelets are approximately elliptical in 2D models, while activated platelets are approximately circular. Piecewise linear approximations of platelets are currently used in IB methods applied to the simulation of platelet aggregation (*e.g.,* [4, 13, 14]).

In previous work [17], we found that interpolation with RBFs restricted to the circle (or sphere in 3D) offered accuracy and computational cost comparable with that offered by Fourier-based methods in modeling an infinitely smooth target shape, its normals and tension forces computed on its surface. Furthermore, interpolation with RBFs resulted in better convergence (often an order more) than that offered by both Fourier-based methods when the target shape had only one or two underlying derivatives. In general, use of RBFs led to a computational cost comparable with that of Fourier-based methods and orders of magnitude lower (for the same accuracy) than that of the standard combination of techniques (piecewise quadratic interpolation and finite differences) used in many IB methods. This RBF-based geometric model has since been used in a variety of applications [18–20]. It has recently been extended to the representation of open elastic curves immersed in a purely viscous fluid for use within the method of regularized Stokeslets [21].

We now turn our attention to exploring the consequences of using this parametric RBF geometric model within the full IB method, with platelet aggregation as our target application. We seek to determine if the advantages inherent in the RBF interpolation of static shapes carry over to full-fledged IB simulations, and if the RBF interpolation can give us benefits that are apparent only in full-fledged IB simulations. In this work, we propose a new IB algorithm that utilizes the features afforded by our RBF geometric model.

The paper is organized as follows. In Section 2, we briefly discuss the traditional IB method for simulating fluid–structure interaction. In Section 3, we review not only the piecewise linear and RBF geometric modeling strategies but also the components necessary for handling immersed elastic structures in the IB method. In Section 4, we provide details of the spatial and temporal discretizations of both versions of the IB method. In Section 5, we present our comparison of the RBF-IB method with the traditional IB method in terms of convergence, accuracy, area loss, and time-step size. We also provide energy estimates for RBF-IB simulations. We then present results from a large platelet aggregation simulation in 2D. Section 6 contains a summary of our findings and a discussion of future research directions.

Notation: We denote vectors with as many components as the spatial dimension in bold. We denote vectors with as many components as the number of data sites (N_{d}) or sample sites (N_{s}) by underlining. We indicate matrices with (N_{d}) or (N_{s}) rows and two columns in bold with underlining.

2. REVIEW OF THE IMMERSED BOUNDARY METHOD

To review the IB method, we focus on a simple 2D model problem in which a single fluid-filled closed elastic membrane is immersed in a viscous fluid. The physics of the model problem is that an elastic membrane is under tension and exerts forces on the adjacent fluid. These forces may cause the fluid to move and, correspondingly, cause points on the membrane to move along with the fluid. In the IB method, the fluid is described in the Eulerian frame through a velocity field $\boldsymbol{u}(\boldsymbol{x}, t)$ and pressure field $p(\boldsymbol{x}, t)$ defined at every point $\boldsymbol{x}$ in the physical domain Ω. The elastic membrane is described in the Lagrangian frame. Let the elastic membrane be parameterized by $q \in \Gamma$ and denote by $\boldsymbol{X}(q, t)$ the spatial coordinates at time t of the membrane point labeled by q. The IB equations are the following coupled equations of motion for the fluid variables $\boldsymbol{u}(\boldsymbol{x}, t)$ and $p(\boldsymbol{x}, t)$ and the membrane configuration $\boldsymbol{X}(q, t)$.

$$\rho(\boldsymbol{u}_t + \boldsymbol{u} \cdot \nabla \boldsymbol{u}) = -\nabla p + \mu \nabla^2 \boldsymbol{u} + \boldsymbol{f}, \qquad \nabla \cdot \boldsymbol{u} = 0, \tag{1}$$

$$\boldsymbol{F}(q, t) = \boldsymbol{F}\left(\boldsymbol{X}(q, t), \frac{\partial}{\partial q} \boldsymbol{X}(q, t)\right), \tag{2}$$

$$\boldsymbol{f}(\boldsymbol{x}, t) = \int_{\Gamma} \boldsymbol{F}(q, t) \delta(\boldsymbol{x} - \boldsymbol{X}(q, t)) \, dq, \tag{3}$$

$$\frac{\partial \boldsymbol{X}}{\partial t}(q, t) = \int_{\Omega} \boldsymbol{u}(\boldsymbol{x}, t) \, \delta(\boldsymbol{x} - \boldsymbol{X}(q, t)) d\boldsymbol{x}. \tag{4}$$

Equation (1) is the Navier–Stokes equations that describe the dynamics of a viscous incompressible fluid of constant density ρ and constant viscosity μ, driven by a force density $\boldsymbol{f}$, which here arises because of the elastic deformation of the immersed membrane. Equation (2) specifies the elastic force (per unit q) at each point of the IB object. The functional dependence of this force on the state of the boundary is specified appropriately to the material being modeled. Equation (3) defines the fluid force density $\boldsymbol{f}(\boldsymbol{x}, t)$ in terms of the IB elastic force density $\boldsymbol{F}$. Equation (4) specifies that the velocity of each IB point equals the fluid velocity at the same location, a formulation of the no-slip boundary condition for viscous flows. In the model problem and the platelet applications, we assume that the IB objects are neutrally buoyant; the IB membrane itself carries no mass, and each object's mass is attributed to the fluid in which it sits; for more on the IB method, see [22].

3. GEOMETRIC MODELING OF PLATELETS

In this section, we review the two geometric modeling strategies to be compared in the context of the IB method applied to platelet aggregation; for a full description of these strategies, see [17].

3.1. Piecewise linear model

In the traditional IB method, the surfaces of platelets are represented by a collection of IB points. We henceforth refer to the IB points in the traditional IB method as *sample sites* and denote them by $X_s(t) = X(q, t)$ for each discrete $q \in \Gamma$ and at a particular time t. The surface elastic forces of the platelets are spread from these sample sites into the neighboring fluid. Both tension and bending forces are computed using a finite difference discretization of force models at sample sites. An explicit piecewise linear interpolant of the surface is not formed. If other information (such as normal vectors) is needed at the sample points, an approximation to the surface represented by the sample sites may be formed by a piecewise quadratic interpolation of the sample sites (*e.g.*, [23]). After the incompressible Navier Stokes equations are solved, velocities from the portions of the Eulerian grid surrounding the sample sites are interpolated to the sample sites using a discretization of Equation (4) and used to move the platelets.

3.2. Parametric RBF model

The RBF method is a popular tool for approximating multidimensional scattered data; for an overview of the theory and application of this method, see the books by Fasshauer [24] and Wendland [25]. The restriction of the RBF method to interpolation on a circle and/or sphere is discussed by Fasshauer and Schumaker [26, §6]. When restricted to these domains, the RBF method is referred to as the *spherical basis function* method [25, Ch. 17]. Several studies have provided error estimates for RBF interpolation on circles and spheres; in fact, these interpolants can provide spectral accuracy, provided the underlying target function is sufficiently smooth [27, 28]. The RBF method has also been used successfully for numerically solving partial differential equations on the surface of the sphere [29, 30], as well as more general surfaces [20, 31].

Here, we present the RBF model developed in our earlier work [17]. It is based on explicit parametric representations of the objects. Because our target objects are platelets, which in 2D models are nearly elliptical or circular, we choose a polar parameterization. We use our model to define operators necessary for the computation of geometric and mechanical quantities required by the IB method.

We represent a platelet surface at any time t parametrically by

$$\mathbf{X}(\lambda, t) = (X(\lambda, t), Y(\lambda, t)) \tag{5}$$

where $0 \leqslant \lambda \leqslant 2\pi$ is the parametric variable and $\mathbf{X}(0, t) = \mathbf{X}(2\pi, t)$. We explicitly track a finite set of N_d points $\mathbf{X}_1^d(t), \ldots, \mathbf{X}_{N_d}^d(t)$, which we refer to as *data sites*. Here, $\mathbf{X}_j^d(t) := \mathbf{X}\left(\lambda_j^d, t\right)$, $j = 1, \ldots, N_d$, and we refer to the parametric coordinates $\lambda_1^d, \ldots, \lambda_{N_d}^d$ as the *data site nodes* (or simply *nodes*). We construct each component of $\mathbf{X}$ by using a smooth RBF interpolant of the data sites in parameter space as discussed in detail in the following. We also make use of derivatives of the interpolant at the data sites, and we use the interpolant and its derivatives at another set of prescribed *sample points* or *sample sites*, which correspond to N_s parameter values: $\lambda_1^s, \ldots, \lambda_{N_s}^s$.

We first explain how to construct an RBF interpolant to the X component of $\mathbf{X}$ using the data $\left(\lambda_1^d, X_1^d(t)\right), \ldots, \left(\lambda_{N_d}^d, X_{N_d}^d(t)\right)$; the construction of the Y component follows in a similar manner. Let $\phi(r)$ be a scalar-valued radial kernel, whose choice we discuss in the following. Define $X(\lambda, t)$ by

$$X(\lambda, t) = \sum_{k=1}^{N_d} c_k^X \phi\left(\sqrt{2 - 2\cos\left(\lambda - \lambda_k^d\right)}\right). \tag{6}$$

Note that the square root term in Equation (6) is the Euclidean distance between the points on the unit circle whose angular coordinates are λ and λ_k^d. We have found that the distance argument $r = \sqrt{2 - 2\cos\left(\lambda - \lambda_k^d\right)}$ is far more accurate for the geometric of modeling static closed curves and surfaces than, say, $r = |\lambda - \lambda_k^d|$. In addition, recent work has shown that the periodic distance argument gives results identical to those given by the non-periodic distance argument for

the modeling of both static *and* dynamic open curves [21]. While other distance arguments could be considered (for example, using geodesic distance in place of Euclidean distance), Fuselier and Wright have shown that RBF interpolation can produce favorable error estimates in the interpolation of functions on submanifolds of $\mathbb{R}^n$ even when no knowledge of the surface metric is used [32]. For these reasons, we restrict our attention to the periodic distance argument in (6). For this paper, we use the multiquadric (MQ) radial kernel function, given by

$$\text{MQ:} \quad \phi(r) = \sqrt{1 + (\varepsilon r)^2}, \tag{7}$$

where ε is called the shape parameter. The choice of ε is discussed in Section 5. To have $X(\lambda, t)$ interpolate the given data, we require that the coefficients $c_k^X, k = 1, ..., N_\text{d}$ satisfy the following system of equations:

$$\underbrace{\begin{bmatrix} \phi(r_{1,1}) & \cdots & \phi(r_{1,N_\text{d}}) \\ \phi(r_{2,1}) & \cdots & \phi(r_{2,N_\text{d}}) \\ \vdots & \ddots & \vdots \\ \phi(r_{N_\text{d},1}) & \cdots & \phi(r_{N_\text{d},N_\text{d}}) \end{bmatrix}}_{A} \underbrace{\begin{bmatrix} c_1^X \\ c_2^X \\ \vdots \\ c_{N_\text{d}}^X \end{bmatrix}}_{\underline{c}_\text{d}^X} = \underbrace{\begin{bmatrix} X_1^\text{d}(t) \\ X_2^\text{d}(t) \\ \vdots \\ X_{N_\text{d}}^\text{d}(t) \end{bmatrix}}_{\underline{X}_\text{d}(t)}, \tag{8}$$

where $r_{j,k} = \sqrt{2 - 2\cos\left(\lambda_j^\text{d} - \lambda_k^\text{d}\right)}$. Because $r_{j,k} = r_{k,j}$, the matrix A in this system is symmetric. More importantly, for the MQ kernels, A is non-singular, with the global support and infinite smoothness of $\phi(r)$ lending itself to spectral accuracy and convergence on smooth problems [24, 25]. One could alternatively use any of the other infinitely smooth kernels like the Gaussian or the inverse MQ in place of the MQ kernel.

In our application, we want to be able to evaluate $X(\lambda, t)$ at sample sites corresponding to parameter values $\lambda_1^\text{s}, ..., \lambda_{N_\text{s}}^\text{s}$, which stay fixed over time. While we could use Equation (6) to do this, it is much more convenient from a notational and computational perspective to construct an *evaluation matrix* that combines the linear operations of constructing the interpolant to $\underline{X}_\text{d}(t) = [\underline{X}_\text{d}(t), \ \underline{Y}_\text{d}(t)]$, for any t, and evaluating it at $\lambda_1^\text{s}, ..., \lambda_{N_\text{s}}^\text{s}$. The evaluation matrix can be constructed by first noting that Equation (6) can be written as

$$X(\lambda, t) = \underbrace{\left[\phi\left(\sqrt{2 - 2\cos\left(\lambda - \lambda_1^\text{d}\right)}\right) \cdots \phi\left(\sqrt{2 - 2\cos\left(\lambda - \lambda_{N_\text{d}}^\text{d}\right)}\right)\right]}_{\underline{b}(\lambda)^T} \underline{c}_\text{d}^X. \tag{9}$$

Because $\underline{c}_\text{d}^X = A^{-1}\underline{X}_\text{d}(t)$, we can write Equation (6) as $X(\lambda, t) = \underline{b}(\lambda)^T A^{-1}\underline{X}_\text{d}(t)$. The evaluation of $X(\lambda, t)$ at $\lambda_1^\text{s}, ..., \lambda_{N_\text{s}}^\text{s}$ can then be obtained as follows:

$$\underbrace{\begin{bmatrix} X\left(\lambda_1^\text{s}, t\right) \\ \vdots \\ X\left(\lambda_{N_\text{s}}^\text{s}, t\right) \end{bmatrix}}_{\underline{X}_\text{s}(t)} = \underbrace{\begin{bmatrix} \underline{b}\left(\lambda_1^\text{s}\right)^T \\ \vdots \\ \underline{b}\left(\lambda_{N_\text{s}}^\text{s}\right)^T \end{bmatrix}}_{B} A^{-1}\underline{X}_\text{d}(t) = \underbrace{BA^{-1}}_{\mathcal{E}_\text{s}}\underline{X}_\text{d}(t). \tag{10}$$

So, given the data sites $\underline{X}_\text{d}(t)$ at any time t, we can interpolate their coordinates with an RBF expansion *and* evaluate the interpolant at the sample site nodes $\lambda_1^\text{s}, ..., \lambda_{N_\text{s}}^\text{s}$ to get $\underline{X}_\text{s}(t)$ by the matrix-vector product $\mathcal{E}_\text{s}\underline{X}_\text{d}(t)$. In fact, this same procedure can be used to give values at sample site nodes for any quantity whose values we have at data site nodes and which we represent using an RBF expansion (*e.g.*, $\underline{Y}_\text{d}(t) = \left[Y_1^\text{d}(t) \cdots Y_{N_\text{d}}^\text{d}(t)\right]^T$). Furthermore, the evaluation matrix $\mathcal{E}_\text{s}$ can be precomputed once at $t = 0$ and used for all subsequent times.

We also need to compute geometric quantities, such as tangent vectors, and mechanical quantities such as forces at data sites and/or sample sites. These quantities require computing derivatives with respect to λ of the platelet surface coordinates $(X(\lambda, t), Y(\lambda, t))$. We use the RBF-based representation of the surface to compute these derivatives, and we will express derivatives of the RBF interpolant in matrix-vector form. Toward this end, we use similar notation to Equation (9) and define the vector

$$
\begin{aligned}
\underline{b}_\lambda^n(\tilde{\lambda}) &:= \left. \frac{\partial^n}{\partial \lambda^n} \underline{b}(\lambda) \right|_{\lambda = \tilde{\lambda}} \\
&= \left[\left. \frac{\partial^n}{\partial \lambda^n} \phi\left(\sqrt{2 - 2\cos\left(\lambda - \lambda_1^{\mathrm{d}}\right)}\right) \right|_{\lambda = \tilde{\lambda}} \cdots \left. \frac{\partial^n}{\partial \lambda^n} \phi\left(\sqrt{2 - 2\cos\left(\lambda - \lambda_{N_{\mathrm{d}}}^{\mathrm{d}}\right)}\right) \right|_{\lambda = \tilde{\lambda}} \right]^T,
\end{aligned}
$$

for any $0 \leqslant \tilde{\lambda} \leqslant 2\pi$. Just as $\underline{b}(\tilde{\lambda})^T A^{-1} \underline{X}_{\mathrm{d}}(t)$ gives the value of $X(\tilde{\lambda}, t)$, we can use $\underline{b}_\lambda^n(\tilde{\lambda})$ to obtain the n^{th} derivative of $X(\lambda, t)$ with respect to λ as

$$
\left. \frac{\partial^n}{\partial \lambda^n} X(\lambda, t) \right|_{\lambda = \tilde{\lambda}} = \underline{b}_\lambda^n(\tilde{\lambda})^T A^{-1} \underline{X}_{\mathrm{d}}(t).
$$

The evaluation of the n^{th} derivative of $X(\lambda, t)$ at the data site nodes $\lambda_1^{\mathrm{d}}, \ldots, \lambda_{N_{\mathrm{d}}}^{\mathrm{d}}$ can then be obtained as follows:

$$
\begin{bmatrix} \left. \frac{\partial^n}{\partial \lambda^n} X(\lambda, t) \right|_{\lambda = \lambda_1^{\mathrm{d}}} \\ \vdots \\ \left. \frac{\partial^n}{\partial \lambda^n} X(\lambda, t) \right|_{\lambda = \lambda_{N_{\mathrm{d}}}^{\mathrm{d}}} \end{bmatrix} = \underbrace{\begin{bmatrix} \underline{b}_\lambda^n\left(\lambda_1^{\mathrm{d}}\right)^T \\ \vdots \\ \underline{b}_\lambda^n\left(\lambda_{N_{\mathrm{d}}}^{\mathrm{d}}\right)^T \end{bmatrix}}_{B_{\lambda^{\mathrm{d}}}^n} A^{-1} \underline{X}_{\mathrm{d}}(t) = \underbrace{B_{\lambda^{\mathrm{d}}}^n A^{-1}}_{\mathcal{D}_{\lambda^{\mathrm{d}}}^n} \underline{X}_{\mathrm{d}}(t). \tag{11}
$$

In a similar manner, the evaluation of the n^{th} derivative of $X(\lambda, t)$ at the sample site nodes $\lambda_1^{\mathrm{s}}, \ldots, \lambda_{N_{\mathrm{s}}}^{\mathrm{s}}$ can be obtained by

$$
\begin{bmatrix} \left. \frac{\partial^n}{\partial \lambda^n} X(\lambda, t) \right|_{\lambda = \lambda_1^{\mathrm{s}}} \\ \vdots \\ \left. \frac{\partial^n}{\partial \lambda^n} X(\lambda, t) \right|_{\lambda = \lambda_{N_{\mathrm{s}}}^{\mathrm{s}}} \end{bmatrix} = \underbrace{\begin{bmatrix} \underline{b}_\lambda^n\left(\lambda_1^{\mathrm{s}}\right)^T \\ \vdots \\ \underline{b}_\lambda^n\left(\lambda_{N_{\mathrm{s}}}^{\mathrm{s}}\right)^T \end{bmatrix}}_{B_{\lambda^{\mathrm{s}}}^n} A^{-1} \underline{X}_{\mathrm{d}}(t) = \underbrace{B_{\lambda^{\mathrm{s}}}^n A^{-1}}_{\mathcal{D}_{\lambda^{\mathrm{s}}}^n} \underline{X}_{\mathrm{d}}(t). \tag{12}
$$

For given data sites $\underline{X}_{\mathrm{d}}(t)$ at any time t, we can interpolate these values with an RBF expansion *and* evaluate the n^{th} derivative of the interpolant at the data site nodes by the matrix–vector product $\mathcal{D}_{\lambda^{\mathrm{d}}}^n \underline{X}_{\mathrm{d}}(t)$ and at the sample site nodes by $\mathcal{D}_{\lambda^{\mathrm{s}}}^n \underline{X}_{\mathrm{d}}(t)$. We refer to the $N_{\mathrm{d}} \times N_{\mathrm{d}}$ matrices $\mathcal{D}_{\lambda^{\mathrm{d}}}^n$ and the $N_{\mathrm{s}} \times N_{\mathrm{d}}$ matrices $\mathcal{D}_{\lambda^{\mathrm{s}}}^n$ as *RBF differentiation matrices*.

The matrices $\mathcal{D}_{\lambda^{\mathrm{d}}}^n$ and $\mathcal{D}_{\lambda^{\mathrm{s}}}^n$ can be used to give values at respective data site or sample site nodes of the n^{th} derivative of the RBF interpolant of any quantity whose values we have at the data site nodes (*e.g.*, $\underline{Y}_{\mathrm{d}}(t) = \left[Y_1^{\mathrm{d}}(t) \cdots Y_{N_{\mathrm{d}}}^{\mathrm{d}}(t) \right]^T$). These matrices can also be precomputed once at $t = 0$ and used for all subsequent times.

Having defined the operators to compute derivatives of the RBF interpolant, we define the quantity

$$
\boldsymbol{\tau} := \frac{\partial}{\partial \lambda} \boldsymbol{X}(\lambda, t) = \left(\frac{\partial}{\partial \lambda} X(\lambda, t), \frac{\partial}{\partial \lambda} Y(\lambda, t) \right) = (\tau_X, \tau_Y). \tag{13}
$$

The unit tangent vector to $X(\lambda, t)$ is then given by

$$\hat{\boldsymbol{\tau}} := \frac{\boldsymbol{\tau}}{\|\boldsymbol{\tau}\|} = (\hat{\tau}_X, \hat{\tau}_Y). \tag{14}$$

In our experiments, we assume that the Lagrangian force at a point on a platelet is the sum of a tension force, a bending-resistant force and possibly a force due to a bond between that point and a point on another platelet or the vessel wall. For the tension force, we use the fiber model defined in [22], according to which the elastic tension force density at $X(\lambda_i, t_k)$ is given by

$$\boldsymbol{F}^{\mathrm{T}}(\lambda_i, t_k) = \left. \frac{\partial}{\partial \lambda}(T \hat{\boldsymbol{\tau}}) \right|_{\lambda_i, t_k}, \tag{15}$$

where $T = k_{\mathrm{t}}(\|\boldsymbol{\tau}\| - l_0)$ is the fiber tension and $k_{\mathrm{t}} > 0$ is constant. We set $l_{0,i} = \|\boldsymbol{\tau}\|_{\lambda_i, t_0}$, where t_0 is the initial time of the simulation. For a bending-resistant force, we use a linear variant of the force defined in [33] and define the elastic force density at $X(\lambda_i, t_k)$ because of how much the platelet surface there is bent to be

$$\boldsymbol{F}^{\mathrm{B}}(\lambda_i, t_k) = -k_{\mathrm{b}} \left(\frac{\partial^4 X}{\partial \lambda^4} - \frac{\partial^4 X^0}{\partial \lambda^4} \right) \Bigg|_{\lambda_i, t_k}. \tag{16}$$

Here, $X^0 = X(\lambda_i, t_0)$ is the initial configuration of the platelet and $k_{\mathrm{b}} > 0$ is constant. Ideally, the constants k_{t} and k_{b} would be chosen to reflect values determined from experiments involving real platelets. In our work, we choose k_{t} and k_{b}, which keep isolated platelets in simple shear flows approximately rigid, and scale them as we refine the background Eulerian grid; this scaling ensures that the mechanical properties of the elastic material converge as the background grid is refined [34]. Despite this approximate rigidity of isolated platelets, these platelets may deform significantly because of interactions (binding and unbinding) with other platelets, and in other models, as a consequence of platelet activation also.

We defer discussion of how we compute the forces given by Equations (15) and (16) to the next section (and the Appendix), because the implementation is different for the RBF and piecewise-linear representations of the platelet boundary. However, the force acting on a platelet because of other platelets (and/or walls) is common to both methods. We use the spring force defined in [14]: let $p_1, p_2, ..., p_{N_p}$ be the indices corresponding to the platelets in the domain. Let p_1 and p_2 be the indices of two platelets, which are linked at sample sites $X_{p_1}\left(\lambda_{i_1}^s\right)$ and $X_{p_2}\left(\lambda_{i_2}^s\right)$. The force at $X_{p_1}\left(\lambda_{i_1}^s\right)$ is then given by:

$$\boldsymbol{F}_{p_1}^{\mathrm{C}}\left(\lambda_{i_1}^s, t_k\right) = K_{\mathrm{C}}\left(\|X_{p_2}\left(\lambda_{i_2}^s\right) - X_{p_1}\left(\lambda_{i_1}^s\right)\| - l_{0,\mathrm{C}}\right) \frac{X_{p_2}\left(\lambda_{i_2}^s\right) - X_{p_1}\left(\lambda_{i_1}^s\right)}{\|X_{p_2}\left(\lambda_{i_2}^s\right) - X_{p_1}\left(\lambda_{i_1}^s\right)\|}, \tag{17}$$

where K_{C} and $l_{0,\mathrm{C}}$ are the interplatelet cohesion spring stiffness and the resting length, respectively; we also set $\boldsymbol{F}_{p_2}^{\mathrm{C}}\left(\lambda_{i_2}^s, t_k\right) = -\boldsymbol{F}_{p_1}^{\mathrm{C}}\left(\lambda_{i_1}^s, t_k\right)$. The formulation for platelet–wall links is similar.

4. NUMERICAL DISCRETIZATION

In this section, we present the implementation details for both IB methods. For each method, we briefly describe the spatial discretization for both the Lagrangian and Eulerian quantities. We then describe the time-stepping scheme for each method.

4.1. The piecewise-linear IB method

Traditionally, finite-difference approximations of Equations (15) and (16) are used in conjunction with piecewise linear methods in 2D (*e.g.*, [33]). We use a second-order central difference involving sets of sample sites or IB points to discretize the derivatives involved in the computation of both the

tension and bending forces (including tangent lengths). It is useful to think of these finite-difference approximations to the constitutive model as Hookean springs connecting pairs of IB points. Note that these differences are only second-order assuming a near-uniform sampling. This is one of the sources of error for the IB method.

For the Eulerian spatial discretization, we use a second-order centered finite-difference approximation to the Laplacian on a staggered Marker-and-Cell (MAC) grid [35]. We discretize the advection term (in conservative form $\nabla \cdot (\boldsymbol{u}\boldsymbol{u}^T)$) using second-order centered differences, averaging quantities to cell edges or nodes as required. For the approximate δ-function, we use the 'cosine' form described by Peskin [22], which ensures that the entire IB force is transmitted to the grid, that the force density on the grid is a continuous function of the IB point locations, and that the communication between grid and IB points is very localized. We note that one could also use other discrete δ-functions designed to satisfy specific properties, if required [36]. To prevent leakage, the tension stiffness k_t is set sufficiently high so that the IB point spacing on the surface is approximately $0.5h$ (or less), where h is the Eulerian grid cell width. After each update of the IB point locations, new links are formed and existing ones are broken using the model's rules for these types of events.

We use the formally second-order Runge–Kutta time-stepping scheme outlined in [37]. This time-stepping scheme demonstrates second-order convergence in time for a smooth-forcing function, or for an elastic material that fills the whole domain, as demonstrated in [37]. This scheme exhibits only first-order convergence in time in the presence of a sharp interface between the fluid and the elastic material, as is typical of IB methods. The full scheme is presented in Appendix B.

4.2. The RBF–IB method

In order to construct the operators utilized by our algorithm, we must first choose an appropriate node set. We use N_d equally spaced values on the interval $(0, 2\pi]$ as the data site node set $\left\{\lambda_k^d\right\}_{k=1}^{N_d}$. This gives a uniform sampling in the parametric space. We also use $N_s > N_d$ (typically, $N_s = 4N_d$ or $N_s = 8N_d$) equally spaced points in the interval $(0, 2\pi]$ as the set of sample site nodes $\left\{\lambda_j^s\right\}_{j=1}^{N_s}$ because this results in a set of sample sites that are well distributed over the object. As in the traditional IB method, we make sure to start simulations with a sample site spacing of less than $0.5h$ (again enforced approximately using the tension stiffness), although the data site spacing can be much greater. In the Results section, we explore the ramifications of this choice.

We have formulated our operators to ensure that operations like evaluation of the interpolant and computing derivatives (and therefore the constitutive model) do not require solving a linear system for any time-step of the platelet simulation except the initial step. This is possible because although the data sites and sample sites move over the course of the simulation, their values in parameter space do not change. For the RBF model of the platelets, the evaluation matrix $\mathcal{E}_s$ in Equation (10) and differentiation matrices $\mathcal{D}_{\lambda^d}^n$ and $\mathcal{D}_{\lambda^s}^n$ in Equations (11) and (12), respectively, can be computed using the FFT as discussed in our previous work [17]. This is possible because the data site nodes $\left\{\lambda_k^d\right\}_{k=1}^{N_d}$ are equally spaced, which results in the A matrix defined Equation (8) having a circulant matrix structure. The costs and accuracy of the RBF models are elaborated upon in the discussion of the results. The algorithm to compute forces on platelets using these operators is presented in Appendix A.

The RBF-IB method uses the same time-stepping scheme and Eulerian discretization as the piecewise linear IB method, with one important difference. When computing the forces at time level $n + 1/2$, we advance the *data sites* to time level $n + 1/2$, generate a set of sample sites at that time level, and compute forces at the sample sites. Similarly, we use the mid-step approximation to the velocity field to advance the *data sites* to time level $n + 1$. We thus generate only a single set of sample sites every time-step, because the sample sites are only needed when the data sites are advanced to time level $n + 1/2$. It is clear that if the number of data sites is fewer than the number of sample sites, this results in improved computational efficiency over the piecewise linear IB method. However, it is important to explore the effect of our changes on the convergence of the algorithm. We explore these questions in the Results section; for a more complete description of the RBF-IB time-stepping scheme, see Appendix C.

5. RESULTS

In this section, we first compare the convergence of the RBF-IB method on a canonical test problem. We also use this test problem to explore the relationship between the number of data sites (N_d) and the Eulerian grid spacing (h). We then compare the area loss in an elastic object simulated by each method on the same problem and discuss the time-step sizes allowed by both methods. We follow with a discussion of the change in energy over time in the RBF-IB method. We then provide timings for platelet simulations and discuss both foreseen and unforeseen results of using the RBF model within the IB method. Finally, we present the results of platelet aggregation simulations conducted using the RBF-IB method.

Description of our standard fluid–structure interaction problem:

We describe a standard fluid–structure interaction problem on which we test both versions of the IB method. This problem is commonly used in the IB literature (*e.g.*, [38]). The problem involves placing an elliptical object with its center of mass at the center of the $[0, 1]^2$ physical domain. The elliptical object has a circle of the same area as its rest configuration and attempts to attain the rest configuration subject to a combination of tension and bending forces. The physical domain is filled with a fluid that is initially at rest, with periodic boundary conditions in the x-direction and no-slip Dirichlet boundary conditions in the y-direction. We set the radius of the target circle to be $r = 0.1$ units, with the ellipse initially having a major axis of $a = 2r$ and a minor axis of $b = 0.5r$. This test is visualized in Figure 1.

5.1. Convergence studies

In previous work [17], we compared the accuracy and convergence of both RBF and traditional IB geometric modeling strategies for static platelet-like shapes. We now compare the accuracy and convergence of the full RBF-IB and piecewise-linear IB (PL-IB) methods, both for the velocity field and for the immersed elastic structure.
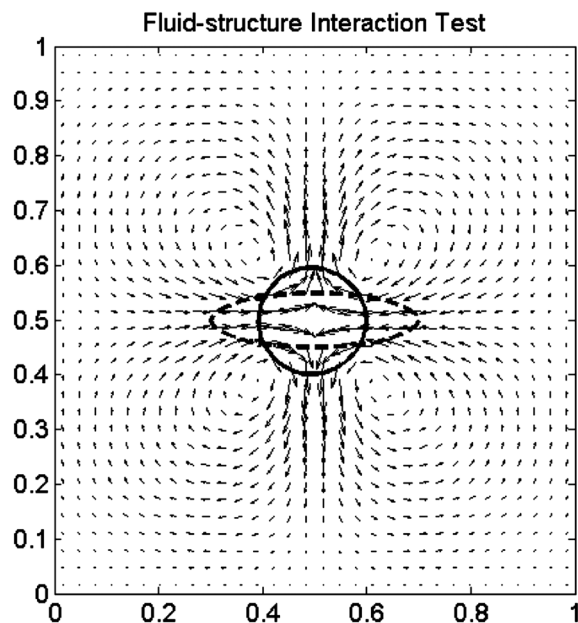

Figure 1. A visualization of the fluid–structure interaction test. The dashed lines show the initial ellipse, while the filled line indicates the near-circular object at the final time $t = 2.0$. The small arrows indicate the velocity field at the final time. The maximum velocity is very close to zero at this time as the object is almost at rest.

For the fluid, on each grid with cell width h, we define the quantity $\boldsymbol{u}^{c,h}$, the coarsened discrete velocity field from the 256×256 grid. This coarsened velocity is obtained using a spline interpolation of each component of the velocity on the finest grid, and evaluating it at the edges on the coarser grid. For each grid point i on a grid with cell width h, we compute the quantity $f_i^h = ||\boldsymbol{u}_i^h - \boldsymbol{u}_i^{c,h}||$. We define the l_2 error in the velocity field as $e_2(h) = \sqrt{\sum_i \left(f_i^h h\right)^2}$ and the l_∞ error in the velocity field to be $e_\infty(h) = \max_i f_i^h$. The convergence rate for errors $e(2h)$ and $e(h)$ is measured as $p_u = \log_2\left(\frac{e(2h)}{e(h)}\right)$.

We note that convergence rates of numerical solutions can be computed in an alternate manner to the one outlined earlier. Solutions on three grids $\boldsymbol{u}^{4h}$, $\boldsymbol{u}^{2h}$, and $\boldsymbol{u}^h$ are chosen; the errors $e(2h)$ and $e(h)$ are computed by comparing the $4h$ grid with the $2h$ grid, and the $2h$ grid to the h grid, respectively. The rate of convergence is then once again estimated as $p_u = \log_2\left(\frac{e(2h)}{e(h)}\right)$. While the results shown in this section do not employ this method, we have verified that the orders of convergence computed using this method match those presented in this section.

For the Lagrangian markers (sample sites or IB points), we adopt the following procedure:

1. Given the number of sample sites N_s and the radius of the target circle r, we define $\theta = \frac{2\pi}{N_s}$, the angle subtended at the center of the circle if the points were evenly spaced.
2. We define the quantity $C = 2r\sin(0.5\theta)$, the chord length between any two points in a set of evenly spaced points on an ideal circle. We also define C_{exact} to be the ideal chord length for $N_s = 400$.
3. We compute the actual distances d_i between the sample sites (or IB points) for a simulation computed on the 256×256 grid with $N_s = 400$. We then compute the quantities $s_\infty = \max_i |d_i - C_{exact}|$ and $s_2 = (1/N_s)\sqrt{\sum_i |d_i - C_{exact}|^2}$.
4. We compute $s_\infty^{N_s}$ and $s_2^{N_s}$ for $N_s = 50, 100, 200$. We then define the l_2 error to be $e_2(N_s) = |s_2^{N_s} - s_2|$ and the l_∞ error to be $e_\infty(N_s) = |s_\infty^{N_s} - s_\infty|$.

We define the convergence rate for errors $e(N_s)$ and $e(2N_s)$ to be $p_X = \log_2\left(\frac{e(N_s)}{e(2N_s)}\right)$.

By defining the Lagrangian errors in this fashion, we circumvent the fact that the 'correct' sample site spacing is unknown; our aforementioned definitions measure 'errors' against equally spaced points on the circle. However, as can be seen in the following text, these errors converge at a first-order rate (asymptotically). For full transparency, we also record the deviations of the sample sites from an equispaced set of IB points on the finest Lagrangian 'grid'; these are given by the quantities s_∞ and s_2.

Table I. Results of a refinement study with the radial basis function–immersed boundary method with $N_d = 25$ data sites. We show the convergence of the velocity field, with errors measured against the velocity field of a simulation on a 256×256 grid with $N_s = 400$ sample sites, $N_d = 100$ data sites, and $\Delta t = 2.5 \times 10^{-5}$.

Grid size	N_s	Δt	L_2 error	Order of convergence	L_∞ error	Order of convergence
32×32	50	2×10^{-4}	1.7666e-04		8.8682e-04	
64×64	100	1×10^{-4}	1.5097e-04	0.23	5.4255e-04	0.71
128×128	200	5×10^{-5}	8.4247e-05	0.84	3.0738e-04	0.82

Table II. Results of a refinement study with the radial basis function–immersed boundary method with $N_d = 25$ data sites. We show the convergence in the sample site positions, with errors measured against the sample site positions of a simulation on a 256×256 grid with $N_s = 400$ sample sites, $N_d = 100$ data sites, and $\Delta t = 2.5 \times 10^{-5}$.

N_s	Grid size	Δt	L_2 error	Order of convergence	L_∞ error	Order of convergence
50	32×32	2×10^{-4}	3.1188e-06		2.3238e-05	
100	64×64	1×10^{-4}	3.5898e-07	3.12	3.0048e-06	2.95
200	128×128	5×10^{-5}	1.5310e-07	1.23	1.7905e-06	0.75

5.1.1. Convergence of the RBF-IB method for $N_d = 25$. Here, we test the convergence of the RBF-IB method on the fluid–structure interaction problem described earlier. We compare the velocity field and sample site positions with those computed for the same test problem on a 256×256 grid with $N_s = 400$ sample sites and $N_d = 100$ data sites. Tables I and II show the results obtained with $N_d = 25$ data sites for the fluid and the object, respectively.

Table I shows that the errors produced in the RBF-IB method. We see that the RBF-IB method for $N_d = 25$ shows first-order convergence asymptotically. Table II shows results for the errors on the object. The errors on the coarsest grid are high, leading to a higher-than-expected convergence rate in both norms when we measure the errors on a 64×64 grid. Once again, the convergence rate remains close to first-order, as expected. For completion, we note that $s_2 = 4.3694e - 08$ and $s_\infty = 1.1113e - 06$ for the structure on the 256×256 grid for $N_d = 25$ data sites.

5.1.2. Convergence of the RBF-IB method for $N_d = 50$. We repeat the previous test problem with $N_d = 50$ data sites. As before, we compare the velocity field and sample site positions with those computed for the same test problem on a 256×256 grid with $N_s = 400$ sample sites and $N_d = 100$ data sites. Tables III and IV show the results for the fluid and the object, respectively.

Examining Table III, we see that when moving to the finest grid, we have once again recovered first-order convergence. The errors on the 128×128 grid for $N_d = 50$ with the RBF-IB method are close to those on the same grid with $N_d = 25$. Table IV shows errors similar to those seen in Table II, albeit with less erratic convergence. Indeed, we recover first-order convergence in the l_2 norm and close to second-order convergence in the l_∞ norm.

We note that using $N_d = 50$ data sites does not result in significantly better convergence on the structure than $N_d = 25$. There are two possible explanations. The first is that the function representing the shape of the object is of limited smoothness (as seen in our previous work [17]), with higher values of N_d causing the interpolation error to saturate or even increase. The alternate (and more likely) explanation is that, because our RBFs are parameterized on the circle, $N_d = 25$ would already have a very high accuracy when the object becomes a circle, considering the spectral accuracy of RBF interpolation on the circle; in such a scenario, using $N_d = 50$ data sites would only serve to increase the rounding errors in the representation of the structure. The values of s_2 and s_∞ for the structure on the 256×256 grid for $N_d = 50$ data sites are the same as those for $N_d = 25$ data sites.

Table III. Results of a refinement study with the radial basis function–immersed boundary method with $N_d = 50$ data sites. We show the convergence of the velocity field, with errors measured against the velocity field of a simulation on a 256×256 grid with $N_s = 400$ sample sites, $N_d = 100$ data sites, and $\Delta t = 2.5 \times 10^{-5}$.

Grid size	N_s	Δt	L_2 error	Order of convergence	L_∞ error	Order of convergence
32×32	50	2×10^{-4}	5.5617e-03		3.7404e-02	
64×64	100	1×10^{-4}	2.2436e-04	4.63	1.2403e-03	4.91
128×128	200	5×10^{-5}	7.8934e-05	1.51	2.8859e-04	2.10

Table IV. Results of a refinement study with the radial basis function–immersed boundary method with $N_d = 50$ data sites. We show the convergence in the sample site positions, with errors measured against the sample site positions of a simulation on a 256×256 grid with $N_s = 400$ sample sites, $N_d = 100$ data sites, and $\Delta t = 2.5 \times 10^{-5}$.

N_s	Grid size	Δt	L_2 error	Order of convergence	L_∞ error	Order of convergence
50	32×32	2×10^{-4}	3.6212e-05		5.2794e-04	
100	64×64	1×10^{-4}	3.8824e-07	6.54	7.8472e-06	6.07
200	128×128	5×10^{-5}	1.7236e-07	1.17	2.0636e-06	1.93

5.1.3. Convergence of the RBF-IB method for $N_d = 0.25N_s$. In the traditional IB method, the number of IB points depends on the grid spacing h. Typically, the number of IB points is chosen so that the distance between any two sample sites is always less that $0.5h$. In all the earlier tests, we have maintained that relationship for the sample sites in the RBF-IB method. In the RBF-IB method, we always use fewer data sites than sample sites, that is, $N_d < N_s$, with the choice of N_d being justified by the results in previous work [17]. Furthermore, in the earlier tests, we fix N_d even as we refine the fluid grid. For $N_d = 25$, this means that as we refine N_s, the distance between data sites increases from 0.8 to $3.2h$ (at the start of the simulation).

In order to gain intuition on the relationship between N_d and h, we now perform a convergence study (using the same test problem given previously) with increasing values of N_d as h is reduced. To accomplish this, we use values of $N_d = 12, 25, 50, 100$ for $N_s = 50, 100, 200, 400$, that is, we enforce $N_d = 0.25N_s$. We use the solution computed with $N_d = 100$ and $N_s = 400$ on a 256×256 grid as our gold standard, just as we have in all the other tests. Table V shows the results for the fluid. Clearly, the errors are higher and the convergence more erratic than for the fixed $N_d = 50$ tests previously presented, but varying N_d certainly seems to give better convergence than fixing it to $N_d = 25$. However, the convergence in the structure is comparable, with lower errors being achieved than both $N_d = 25$ and $N_d = 50$. This can be seen in Table VI. Unfortunately, the advantages of varying N_d with N_s are not clear. Using $N_d = 50$ yields the lowest errors in the fluid on the finest grid level, and reasonably low errors on the structure for all grid levels. Given the similarity of the errors achieved with $N_d = 50$ to those achieved with increasing N_d, we choose the simpler strategy of using a fixed value of $N_d = 50$ for our tests, although we present timings with $N_d = 25$ as well.

5.1.4. Effect of the shape parameter ε. In previous work [17], we found that the RBF shape parameter $\varepsilon > 0$ had to be selected carefully to achieve spectral accuracy in the representation of the elastic structure. In that work, we found that small values of ε were ideal for interpolating smooth target shapes and larger ones for rougher target shapes. In our tests, we found that the errors depended on ε even in the case of fluid–structure interaction, with smaller values of ε giving the lowest values of s_2^e and s_∞^e on the 256×256 grid. However, as we mentioned in our previous work, lower values of ε can make the RBF interpolation matrix more ill-conditioned. While methods (such as RBF-QR and RBF-RA) have been developed to overcome this poor conditioning [39], they are much more expensive than forming and inverting the standard RBF interpolation matrix. We thus choose a small value of $\varepsilon = 1.2$ for all our tests. When using $N_d = 100$, we use $\varepsilon = 2.0$ (which was verified on a static test case to be accurate to 12 digits). These are the smallest values we were able to pick without the interpolation matrix becoming ill-conditioned, a strategy consistent with the one used in our previous work [17].

Table V. Results of a refinement study with the radial basis function–immersed boundary method with $N_d = 0.25N_s$ data sites. We show the convergence of the velocity field, with errors measured against the velocity field of a simulation on a 256×256 grid with $N_s = 400$ sample sites, $N_d = 100$ data sites, and $\Delta t = 2.5 \times 10^{-5}$.

Grid size	N_s	Δt	L_2 error	Order of convergence	L_∞ error	Order of convergence
32×32	50	2×10^{-4}	4.7909e-04		2.4700e-03	
64×64	100	1×10^{-4}	1.5097e-04	1.67	5.4255e-04	2.19
128×128	200	5×10^{-5}	9.0802e-05	0.73	3.3020e-04	0.72

Table VI. Results of a refinement study with the radial basis function–immersed boundary method. We show the convergence in the sample site positions, with errors measured against the sample site positions of a simulation on a 256×256 grid with $N_s = 400$ sample sites, $N_d = 100$ data sites, and $\Delta t = 2.5 \times 10^{-5}$.

N_s	Grid size	Δt	L_2 error	Order of convergence	L_∞ error	Order of convergence
50	32×32	2×10^{-4}	9.7439e-06		7.0459e-05	
100	64×64	1×10^{-4}	3.5898e-07	4.76	3.0048e-06	4.55
200	128×128	5×10^{-5}	1.2694e-07	1.50	1.4362e-06	1.07

5.2. *Area loss and time-step size in the RBF-IB method:*

In this section, we study the area loss in the RBF-IB method in a refinement study. We then explore the maximum stable time-step size afforded by each IB method.

The PL-IB method generally attempts to maintain an IB point separation distance of $0.5h$ in order to reduce area loss over the coarse of the simulation. In the RBF-IB method, while the sample site spacing is initially set at $0.5h$, we initialize the structure with a much coarser data site discretization, with the data site separation being almost $3.2h$ in some cases. In addition, we use the same strategy for interpolating velocities that we carried out in the PL-IB method, that is, we interpolate velocities to data sites from a 4×4 patch of fluid around each data site. While this can result in significant computational savings, it is important to explore the area loss in our discretization. We turn once again to our standard fluid–structure interaction problem. We run that simulation on successively finer grids until time $t = 2$. For both the RBF-IB method and the PL-IB method, we measure the initial area of the object for the same initial configuration of points. We then measure the area at time $t = 2$ and compute the percentage change in area.

In order to get an accurate estimate of the area in both methods, we fit an RBF interpolant to each object's Lagrangian markers (data sites for the RBF-IB method and all the IB points for the PL-IB method). We then sample that interpolant at a fixed number of points (400 points) and use the trapezoidal rule to compute the area. As was mentioned earlier, we ensure that the initial ellipse has the same area as the target circle by picking its radii to be $a = 2r$ and $b = 0.5r$, where $r = 0.1$ is the radius of the target circle. The exact area is then $\frac{\pi}{100}$. Our approach of sampling each object and computing the area with the trapezoidal rule gives an area estimate that agrees with this value up to seven digits at $t = 0$. We record the results of our refinement study in Table VII.

From the table, it is clear that the area loss for fixed $N_d = 25, 50$ and $N_d = 0.25N_s$ are all close to each other. On the coarsest grid, it appears that smaller values of N_d result in lower area loss. The area losses for $N_d = 50$ match with those given by the PL-IB method (results not shown), except in the case of the coarsest grid, where the PL-IB method gives almost a 1% area loss. The convergence is initially second-order but quickly saturates. This saturation is likely due to two sources of error: the first is the interpolation of velocities to the Lagrangian markers, which does not preserve the divergence-free nature of the fluid velocity; and the second is the fact that the time-integration itself is not specifically designed to preserve area. Nevertheless, it is clear from this study that the RBF-IB method produces similar area losses to the PL-IB method despite using a smaller number of Lagrangian markers to move the structure through the fluid.

Another measure of interest is the maximum stable time-step size afforded by each method. We measure this by increasing the time-step size in small increments and observing the forces produced on the structure in the fluid–structure interaction test. Using a time-step that is too large can result in the forces blowing up and the simulation halting. We immediately note that the PL-IB method allows a maximum time-step size of $\Delta t = 2 \times 10^{-4}$ on the 32×32 grid when $N_s = 50$ IB points are used, and a maximum time-step size of $\Delta t = 10^{-4}$ on the 64×64 grid when $N_s = 100$ IB points are used. We use these values of Δt as the starting point when testing for the time-step sizes allowed by the RBF-IB method, and increase the value of Δt in increments of 10^{-4}. We found that on the 32×32 grid, the RBF-IB method allows us to take time-steps that are $3\times$ larger than the time-steps allowed by the traditional IB method; on the 64×64 grid, the RBF-IB method can use time-steps

Table VII. Percentage area loss in the radial basis function–immersed boundary (RBF-IB) method as a function of grid size, the number of sample sites N_s, and the time-step Δt. The PL-IB method gives area losses similar to the $N_d = 50$ case, except on the coarsest grid, where the percentage area loss is three times that of the RBF-IB method.

N_s	Grid size	Δt	% area loss ($N_d = 25$)	% area loss ($N_d = 50$)	% area loss ($N_d = N_s/4$)
50	32×32	2×10^{-4}	0.0680	0.3081	0.0450
100	64×64	1×10^{-4}	0.0047	0.0049	0.0047
200	128×128	5×10^{-5}	0.0023	0.0025	0.0025
400	256×256	2.5×10^{-5}	0.0015	0.0015	0.0015

that are 1.5× larger than the time-steps allowed by the traditional IB method. This pattern holds both when $N_d = 25$ and $N_d = 50$ data sites are used.

In simulations involving platelet-like shapes (ellipses that attempt to maintain their elliptical configuration), we found that the RBF-IB method allows time-step sizes that are 6× larger than those allowed by the PL-IB method on a 32×32 grid, and 3× larger than those allowed by the PL-IB method on a 64×64 grid. This is likely because of the fact that platelet simulations involve smaller deformations than those seen in the standard fluid–structure interaction test.

5.3. Energy estimates

In this section, we compute energy estimates for the RBF-IB method in the context of our standard fluid–structure interaction problem. We run our simulation out to time $t = 2.0$ on two grid sizes, 32×32 and 64×64, with time-step sizes $\Delta t = 2 \times 10^{-4}$ and $\Delta t = 10^{-4}$, respectively. We use $N_d = 50$ data sites.

In this test, one expects the changes in energy to be mainly due to the deformation of the stiff elastic object. Eventually, the energy of the system must damp out as the elastic object reaches its target configuration. We compute the change in energy to demonstrate that the energy is bounded within the RBF-IB simulation. The energy change in a time-step is computed as the sum of the difference in kinetic energy of the fluid over the time-step and the change in potential energy of the elastic object. This can be written as

$$\Delta E = \sum_{fluid} \rho \boldsymbol{u}^{n+1} \cdot \boldsymbol{u}^{n+1} - \sum_{fluid} \rho \boldsymbol{u}^n \cdot \boldsymbol{u}^n + \Delta t \sum_{\boldsymbol{X}} \boldsymbol{F}^{n+1/2} \cdot \frac{\partial \boldsymbol{X}}{\partial t}^{n+1/2} \tag{18}$$

Here, the Lagrangian force $\boldsymbol{F}$ is computed at time level $n + 1/2$ at the sample sites. The $\frac{\partial \boldsymbol{X}}{\partial t}$ term is computed by applying the evaluation matrix $\mathcal{E}_s$ to the velocities obtained at the data sites. This gives us sample site velocities, allowing us to compute dot products with the $\boldsymbol{F}$ terms.

The results of this test are shown in Figure 2. Both plots show the change in energy of the system for the fluid–structure interaction problem on a 32×32 grid (left) and a 64×64 grid (right). Here, the fluid starts off stationary, so the initial kinetic energy is zero. However, the elliptical elastic


Figure 2. Change in energy *per time-step* as a function of time in the radial basis function–immersed boundary method. The figure on the top left shows the change in energy over a time-step as a function of time for the standard fluid–structure interaction test on a 32×32 grid. The figure on the top right shows the same quantity on a 64×64 grid. We use $N_d = 50$ data sites for both grid sizes. The inset plots show the initial spikes corresponding to the change from an ellipse to a circle, which are difficult to see in the main plots.

object starts off under tension because its target configuration is a circle. This means that the initial elastic potential energy of the system is high (although negative by convention). As the elastic object attempts to minimize its elastic potential energy, its deformation drives a change in the kinetic energy of the fluid, causing the kinetic energy of the fluid to increase from its initial value of zero to some maximum. However, the elastic object soon attains something close to its reference configuration, causing the kinetic energy of the fluid to drop sharply. The spikes in both the left and right sides in Figure 2 correspond to that rise and fall in kinetic energy and the trending of the potential energy to zero on both grids, and can be seen more clearly in the inset plots. The viscosity of the fluid causes the kinetic energy to eventually damp out almost completely, with minor perturbations due to possible deformations of the elastic object. The energy of the system continues to decrease as the object becomes more and more circular. In fact, our estimates show that the change in energy is negative, indicating that our method is dissipative. The results are similar for $N_d = 25$ (not shown), although using more data sites appears to make our method less dissipative on this particular test problem.

5.4. Timings for platelet simulations

We now present timings of simulations involving platelet-like shapes. The setup here is different from the standard fluid–structure interaction test. We place ellipses ($r = 0.05$, $a = 2r$, $b = 0.5r$)
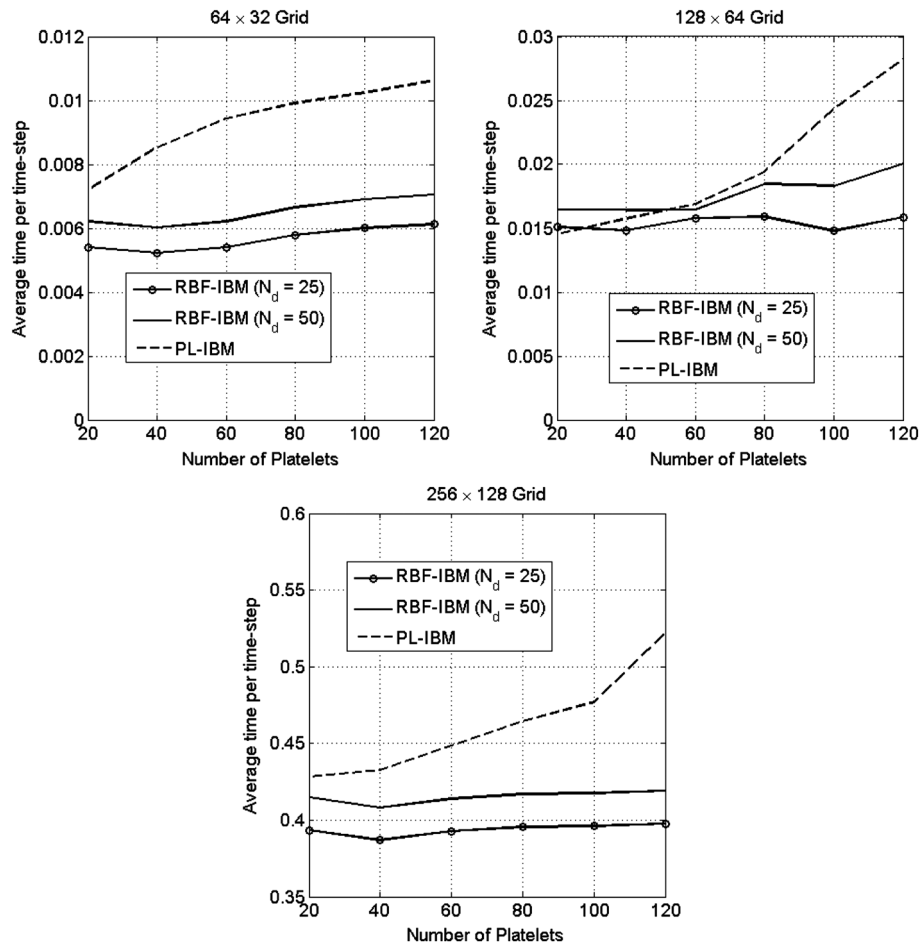

Figure 3. Average time per time-step for 10^5 time-steps of each simulation method as a function of the number of platelets. In the first row, the figure on the left shows timings on a 64×32 grid and the figure on the right on a 128×64 grid. The figure below shows timings on a 256×128 grid. The time-step was set to $\Delta t = 10^{-4}$ for the figures on the top row, and was set to $\Delta t = 10^{-5}$ for the figure on the bottom.

at the left end of a $[0, 2] \times [0, 1]$ domain that resembles a channel. These ellipses represent platelets, and they attempt to maintain their elliptical shapes, that is, their configuration at time $t = 0$ is their preferred configuration. We apply a background force that would result in parabolic velocity field in the absence of the platelets, with a density $\rho = 1.0$ and a non-dimensionalized viscosity of $\mu = 8.0$. The field has a maximum velocity value of $u_{max} = 5.0$, with no-slip boundary conditions on the top and bottom of the domain and periodic boundary conditions at the left and right ends. A platelet is removed from the domain if its center of mass crosses the location $x = 1.9$.

Figure 3 shows timings for three grid sizes for each method as a function of the number of platelets (N_p) being simulated. The number of sample sites was fixed at $N_s = 100$ for both methods, and the number of data sites for the RBF-IB method was set to $N_d = 25$ for one set of tests and then to $N_d = 50$ for the next set. We plot the average time per time-step as a function of the number of platelets; this was computed by running simulations on each grid for 10^5 time-steps, and dividing the total wall-clock time by the number of time-steps. We average this over three runs of each simulation.

While the cost of platelet operations always increases as we increase the number of platelets, the increase in cost is slower (and the absolute cost) for the RBF-IB method because of the RBF representation. For example, for $N_p = 60$, the PL-IB method directly spreads forces from, interpolates to, and moves a total of 6000 IB points (twice per time-step because of the RK2 scheme), while the RBF-IB method with $N_d = 25$ data sites computes forces at 6000 points and interpolates velocities to (and moves) only 1620 points twice per time-step. If the number of platelets is doubled to $N_p = 120$, the PL-IB method now computes forces at, spreads from, interpolates to, and moves 12 000 points twice per time-step, whereas the RBF-IB method computes forces at 12 000 points, but interpolates velocities to and moves only 3240 points twice per time-step. The cost of the RBF-IB method shows better than linear scaling with respect to the number of platelets on all the tested grid sizes for these reasons. Furthermore, it is clear that there is not much of a difference in computational cost between using $N_d = 25$ data sites and $N_d = 50$ data sites.

5.4.1. Effect of the RBF representation on the fluid solver.. In previous work [17], we showed that using an RBF interpolant for geometric modeling is more computationally efficient (for a given accuracy) than using piecewise quadratics and finite differences. However, that benefit alone does not explain the computational efficiency of the RBF-IB method over the PL-IB method that we see in Figure 3.

To fully understand the speedup seen with the RBF-IB method, it is important to understand how the costs are distributed between the different operations (platelet operations and fluid solves) in both IB methods. We show the results for $N_p = 60$ platelets in Table VIII. Clearly, as h is reduced, both IB codes spend more time in the fluid solver than on platelet operations. However, the RBF-IB method clearly spends less time in the fluid solver than the PL-IB method does as we refine the background Eulerian grid.

Indeed, this unexpected result is what gives the RBF-IB method an edge even when the cost of fluid solves dominates the cost of platelet operations. We hypothesize that this may be caused by the RBF representation producing smoother Lagrangian forces than the finite difference model used

Table VIII. Percentage of time per time-step spent in fluid solver as a function of grid size by both methods for $N_p = 60$ platelets. The percentages for the radial basis function–immersed boundary (IB) method are the same for both $N_d = 25$ and $N_d = 50$ data sites, with the total time for the latter being larger. All results use $N_s = 100$ sample sites (or IB points in the PL-IB method) per platelet.

Grid size	% time in fluid solver (RBF-IB)	% time in fluid solver (PL-IB)
64×32	33.7	32.1
128×64	56.2	58.0
256×128	65.4	79.3

in the PL-IB method. Our experiments show that the RBF-IB code needs fewer iterations in the linear solver used in the pressure projection–anywhere from $10 - 30\%$ fewer than the identical fluid solver used in the PL-IB method, depending on the time-step size and the grid resolution, with larger savings on finer grids and smaller time-step sizes.

5.5. Platelet aggregation

We now present the results of a platelet aggregation simulation with the RBF-IB simulation. We used the same boundary conditions, domain size, fluid properties, and Poisseuille flow as in the previous subsection, but allow platelets to form links with other platelets and a portion of the chamber wall ($x = 0.4$ to $x = 0.7$) at the sample sites ($N_s = 100$ per platelet). We used $N_d = 50$ data sites per platelet, making the data sites a subset of the sample sites for convenience of visualization, and then visualize the data sites and the links between sample sites. We allowed each platelet to form up to 10 links in total, either with the wall or with a neighbor; we allow links to cross each other for the purpose of simplicity, although this is usually prohibited in a platelet simulation. The simulation was run on a 128×64 grid with a time-step of $\Delta t = 10^{-4}$.

Each platelet is initially an ellipse with radii $a = 0.06$ and $b = 0.015$. We initialize the platelets so that their centers of masses are at locations $(0.5, 0.02)$, $(0.64, 0.02)$, $(0.78, 0.02)$, $(0.55, 0.07)$, $(0.68, 0.07)$, $(0.4, 0.045)$, $(0.23, 0.045)$, and $(0.65, 0.14)$. We chose these locations to ensure that three platelets lay on the wall, with three close enough to bind to the three bound to the wall, and two slightly further away. Each platelet attempts to maintain its initial elliptical shape. We then started the simulation and ran it to time $t = 2.4$. The results are shown in Figure 4. The figure shows both the velocity field and the platelet aggregate for a portion of $[0, 2] \times [0, 1]$ domain, the data sites on each platelet and the links between the sample sites corresponding to those particular data sites on the platelet.

There are two interesting features in Figure 4. The first is that the fluid flow becomes diverted around the growing aggregate, a consequence of the size of the aggregate and the dynamics of the problem that mimics what one would hope to see in a realistic platelet aggregation simulation. The second feature is that some platelets are quite deformed, for example, the platelet with center of mass approximately at $(0.5, 0.02)$, or its neighbor above and to its left. This is a consequence (and function) of the stiffness of each platelet, the shear rate of the flow, and the number of links we allow each platelet to form. Higher platelet stiffness, lower shear rates, and fewer (or weaker) links would lead to less deformation. The breaking model for inter-platelet and platelet–wall links can also affect the mechanics of aggregation. We note that our RBF model did not run into any instabilities in this simulation even when we ran it out to a time at which all the platelets (except the three closest to the wall) had left the domain.


Figure 4. Results of a platelet aggregation simulations with the radial basis function–immersed boundary (RBF-IB) method. The figure shows a zoomed-in snapshot of a platelet aggregation simulation achieved with the RBF-IB method with a time-step of $\Delta t = 10^{-4}$. The snapshot was taken at simulation time $t = 2.4$. The simulation was run on a 128×64 grid on a $[0, 2] \times [0, 1]$ domain. The arrows show the magnitude and direction of the velocity field.

6. SUMMARY

The IB method, as a numerical methodology for applications involving fluid–structure interactions, naturally lends itself to our problem of interest: simulating platelet aggregation during blood clotting. In this application, platelets are modeled as immersed elastic structures whose shapes change dynamically in response to blood flow and chemistry. In previous work [17], we discussed several geometric representations for platelets and compared them with the representation used within the traditional IB method. We concluded that an RBF geometric model for platelets would prove advantageous in several ways.

In this work, we explored the ramifications of using the RBF geometric model within the IB method and compared the behavior of this new RBF-IB method against that of the traditional IB method. We discussed the issue of selecting an appropriate shape parameter for the RBF-IB method. We then presented a series of convergence studies for measuring errors and convergence both in the velocity field and in the representation of the immersed elastic structure. We went on to compare the computational costs incurred by both methods in the context of platelet simulations. We then compared the area conservation properties of both methods and also the time-step restrictions on both. We also remarked on the energy properties of our method.

We conclude the following:

- The RBF-IB method demonstrates first-order convergence, similar to that shown by the traditional IB method.
- The RBF-IB method can be safely used both with a fixed number of data sites and with an increasing number of data sites as the background Cartesian grid is refined; the correct strategy is likely problem-specific;
- Through the use of a global interpolant and the dual representation (data sites and sample sites), the RBF-IB method allows for a decoupling of accuracy in geometric modeling from accuracy in the full fluid solve, allowing coarse Lagrangian representations when sufficient.
- The RBF-IB method is more computationally efficient than the traditional IB method, both because of the utilization of a small number of data sites and smoother forces being spread into the fluid, resulting in a faster convergence from the fluid solver; and
- The RBF-IB allows for larger time-step sizes than those allowed in the traditional IB method for a given grid size.

In previous work [38], a sufficient condition for unconditional stability of an implicit IB method was established. The proof relied on the assumption that the set of points from which IB forces are spread is the same as that to which grid velocities are interpolated to update IB point positions. The RBF-IB method does not meet that condition, and it remains to be seen how this would impact an implicit version of our method. Finally, an issue with the RBF-IB method is that it is dependent on the parameterization of the immersed elastic objects. For objects that are not easily parameterized in terms of circles and ellipses, the use of the RBF model as presented in our work (wherein the RBFs are restricted to the circle) may not be ideal. In the future, we thus hope to explore the use of RBFs in a meshfree variational form within the IB method so as to be able to easily evaluate constitutive models on arbitrary shapes.

APPENDIX A: ALGORITHM FOR COMPUTING PLATELET FORCES WITH RADIAL BASIS FUNCTIONS

We now describe the implementation of the constitutive models outlined in Section 3. We present algorithms for computing platelet forces in 2D.

Notation: In the description of the following algorithms, we use standard matrix–vector operations such as multiplication and non-standard operations like element-by-element multiplication of matrices and vectors (sometimes called the Hadamard product). We denote this latter operation with the $\circ$ operator. For example, if $\underline{\textbf{\textit{J}}}$ and $\underline{\textbf{\textit{L}}}$ are $N_\mathrm{d} \times 2$ matrices and $\underline{R}$ is a vector of length N_d, then the i^th row of $\underline{\textbf{\textit{J}}} \circ \underline{\textbf{\textit{L}}}$ and $\underline{R} \circ \underline{\textbf{\textit{J}}}$ are given by

$$(\underline{\boldsymbol{J}} \circ \underline{\boldsymbol{L}})_{i,1:2} = [(\underline{\boldsymbol{J}})_{i,1}(\underline{\boldsymbol{L}})_{i,1}, \ (\underline{\boldsymbol{J}})_{i,2}(\underline{\boldsymbol{L}})_{i,2}]$$
$$(\underline{R} \circ \underline{\boldsymbol{L}})_{i,1:2} = [(\underline{R})_{i}(\underline{\boldsymbol{L}})_{i,1}, \ (\underline{R})_{i}(\underline{\boldsymbol{L}})_{i,2}]$$

where $i = 1, \ldots, N_{\mathrm{d}}$.

We define $\underline{\boldsymbol{\tau}}_{\mathrm{d}} = \mathcal{D}^1_{\lambda^{\mathrm{d}}} \underline{\boldsymbol{X}}_{\mathrm{d}}(t)$, the $N_{\mathrm{d}} \times 2$ matrix of tangent vectors at the data sites at time t and $\|\boldsymbol{\tau}_{\mathrm{d}}\|$, the N_{d} vector containing the two-norm of each row of $\underline{\boldsymbol{\tau}}_{\mathrm{d}}$. The algorithm for computing platelet elasticity is as follows:

1. Initialization ($t = t_0$): After creating and storing the RBF evaluation matrix as in Equation (10) and differentiation matrices as in Equations (11) and (12), compute for each platelet:

   (a) The rest lengths for the tension force at the data sites: $\underline{l_0} = \underline{\boldsymbol{\tau}}_{\mathrm{d}} = \mathcal{D}^1_{\lambda^{\mathrm{d}}} \underline{\boldsymbol{X}}_{\mathrm{d}}(t_0)$.
   (b) The bending-resistant force term for the platelet's initial configuration at the data sites, $\mathcal{D}^4_{\lambda^{\mathrm{s}}} \underline{\boldsymbol{X}}_{\mathrm{d}}(t_0)$.

2. For each time-step ($t = t_k, k \geqslant 1$), compute for each platelet:

   (a) The length of the tangent vectors $\underline{\boldsymbol{\tau}}_{\mathrm{d}} = \mathcal{D}^1_{\lambda^{\mathrm{d}}} \underline{\boldsymbol{X}}_{\mathrm{d}}(t_k)$ at the data sites: $\|\boldsymbol{\tau}_{\mathrm{d}}\|$; and the unit tangents at the data sites: $\hat{\underline{\boldsymbol{\tau}}}_{\mathrm{d}}$.
   (b) The tension at the data sites, using the constitutive model: $\underline{T}_{\mathrm{d}} = k_{\mathrm{t}}(\|\boldsymbol{\tau}_{\mathrm{d}}\| - \underline{l_0})$.
   (c) The tension force at sample sites: $\underline{\boldsymbol{F}}^{\mathrm{T}}_{\mathrm{s}} = \mathcal{D}^1_{\lambda^{\mathrm{s}}} \underline{\boldsymbol{Z}}_{\mathrm{d}}$, where $\underline{\boldsymbol{Z}}_{\mathrm{d}} = \underline{T}_{\mathrm{d}} \circ \hat{\underline{\boldsymbol{\tau}}}_{\mathrm{d}}$.
   (d) The bending force at sample sites: $\underline{\boldsymbol{F}}^{\mathrm{B}}_{\mathrm{s}} = -k_{\mathrm{b}} \left(\mathcal{D}^4_{\lambda^{\mathrm{s}}} \underline{\boldsymbol{X}}_{\mathrm{d}}(t_k) - \mathcal{D}^4_{\lambda^{\mathrm{s}}} \underline{\boldsymbol{X}}_{\mathrm{d}}(t_0) \right)$.
   (e) The interplatelet cohesion force from Equation (17) at the sample sites: $\underline{\boldsymbol{F}}^{\mathrm{C}}_{\mathrm{s}}$.
   (f) The total Lagrangian force at the sample sites: $\underline{\boldsymbol{F}}_{\mathrm{s}} = \underline{\boldsymbol{F}}^{\mathrm{T}}_{\mathrm{s}} + \underline{\boldsymbol{F}}^{\mathrm{B}}_{\mathrm{s}} + \underline{\boldsymbol{F}}^{\mathrm{C}}_{\mathrm{s}}$.

APPENDIX B: TIME-STEPPING FOR THE PL-IB METHOD

Here, we present the steps of the traditional IB algorithm when used with the RK2 time-stepping scheme from [37].

1. Advance the structure to time level $t^{n+1/2}$ using the current velocity field on the grid $\boldsymbol{u}^n_g$. This is carried out by updating each IB point $\boldsymbol{X}_q$ (for each q) using the equation

$$X^{n+1/2}_q = X^n_q + \frac{\Delta t}{2} U^n_q \equiv X^n_q + \frac{\Delta t}{2} \sum_g \boldsymbol{u}^n_g \delta_h \left(\boldsymbol{x}_g - \boldsymbol{X}^n_q \right) h^2, \qquad (\mathrm{B.1})$$

where h is the fluid grid spacing and δ_h is a discrete approximation to a 2D δ-function. Here, $\boldsymbol{x}_g$ and $X^{n+1/2}_q$ are the coordinates of grid point g and IB point q, respectively

2. The resultant $\boldsymbol{F}^{n+1/2}_q$ of all of the force contributions that act on an IB point $X^{n+1/2}_q$ is calculated for each q.

3. These forces are distributed to the Eulerian grid used for the fluid dynamics equations using a discrete version of Equation (3):

$$\boldsymbol{f}^{n+1/2}_g \equiv \boldsymbol{f}^{n+1/2}(\boldsymbol{x}_g) = \sum_q \boldsymbol{F}^{n+1/2}_q \delta_h \left(\boldsymbol{x}_g - X^{n+1/2}_q \right) dq. \qquad (\mathrm{B.2})$$

Here, $\boldsymbol{F}^{n+1/2}_q$ is the Lagrangian force (per unit q) on the IB point, dq is the increment in parameter q between consecutive discrete sample sites, and δ_h is the same approximate δ-function as used in Equation (B.1).

4. With the fluid force density $\boldsymbol{f}^{n+1/2}_g$ now known at each grid point, the fluid velocity is updated taking a half step ($\Delta t / 2$) with a discrete Navier–Stokes solver. As in [37], we use a fractional-step projection method. First, a backward Euler discretization of the momentum equations is used. The pressure that enforces discrete incompressibility is determined [40]. This gives us the velocity field $\boldsymbol{u}^{n+1/2}_g$, the mid-step approximation required in an RK2 method.

5. Using the mid-step fluid velocity $\boldsymbol{u}_g^{n+1/2}$ and the mid-step IB point positions $\boldsymbol{X}_q^{n+1/2}$, update the IB points $\boldsymbol{X}_q^n$ for each q to the time level t^{n+1} using

$$\boldsymbol{X}_q^{n+1} = \boldsymbol{X}_q^n + \Delta t\, \boldsymbol{U}_q^{n+1/2} \equiv \boldsymbol{X}_q^n + \Delta t \sum_g \boldsymbol{u}_g^{n+1/2} \delta_h \left(\boldsymbol{x}_g - \boldsymbol{X}_q^{n+1/2} \right) h^2, \qquad \text{(B.3)}$$

where δ_h is the same approximate δ-function we have used throughout.

6. Update the velocity $\boldsymbol{u}_g^n$ to time-level t^{n+1} using the mid-step velocity $\boldsymbol{u}_g^{n+1/2}$ and the force $\boldsymbol{f}_g^{n+1/2}$. The mid-step velocities are advected, while a Crank–Nicolson scheme is used for time-stepping the momentum equations. The pressure projection gives us the discretely incompressible velocity field $\boldsymbol{u}_g^{n+1}$. Note that this step could have been performed as soon as $\boldsymbol{u}_g^{n+1/2}$ was computed. It is independent of step (5).

APPENDIX C: TIME-STEPPING FOR THE RADIAL BASIS FUNCTION–IMMERSED BOUNDARY METHOD

The RBF-IB method is time-stepped using the same RK2 method as earlier, with a few changes to incorporate data sites and sample sites.

1. Advance the structure to time level $t^{n+1/2}$ using the current velocity field $\boldsymbol{u}^n$. This is carried out by updating the *data sites* $(\boldsymbol{X}_d)_j^n$ by a discrete analog of Equation (4)

$$(\boldsymbol{X}_d)_j^{n+1/2} = (\boldsymbol{X}_d)_j^n + \frac{\Delta t}{2}(\boldsymbol{U}_d)_j^n \equiv (\boldsymbol{X}_d)_j^n + \frac{\Delta t}{2} \sum_g \boldsymbol{u}_g^n \delta_h \left(\boldsymbol{x}_g - (\boldsymbol{X}_d)_j^n \right) h^2. \qquad \text{(C.1)}$$

2. Generate a new set of sample sites $\underline{\boldsymbol{X}}_s(t_{n+1/2})$ by applying the RBF evaluation operator to the data sites $\underline{\boldsymbol{X}}_d^{n+1/2} := \underline{\boldsymbol{X}}_d(t_{n+1/2})$, that is,

$$\underline{\boldsymbol{X}}_s(t_{n+1/2}) = \mathcal{E}_s \underline{\boldsymbol{X}}_d(t_{\text{new}}). \qquad \text{(C.2)}$$

3. The total force at the sample sites $\underline{\boldsymbol{F}}_s^{n+1/2}$ is calculated using the algorithm from Appendix A.

4. These forces are distributed to the Eulerian grid used for the fluid dynamics equations using a discrete version of Equation (3):

$$\boldsymbol{f}_g^{n+1/2} \equiv \boldsymbol{f}^{n+1/2}(\boldsymbol{x}_g) = \sum_q \boldsymbol{F}_q^{n+1/2} \delta_h(\boldsymbol{x}_g - (\boldsymbol{X}_s)_q^{n+1/2}) dq. \qquad \text{(C.3)}$$

Here, $\boldsymbol{x}_g$ and $(\boldsymbol{X}_s)_q^{n+1/2}$ are the coordinates of grid point g and sample site q, respectively, $\boldsymbol{F}_q^{n+1/2}$ is the Lagrangian force (per unit q) on the sample site, dq is the increment in parameter q between consecutive discrete sample sites, and δ_h is the same approximate δ-function as used in Equation (C.1).

5. With the fluid force density $\boldsymbol{f}_g^{n+1/2}$ now known at each grid point, the fluid velocity is updated taking a half step $(\Delta t/2)$ with a discrete Navier–Stokes solver. Again, we use a fractional-step projection method, with a backward Euler discretization of the momentum equation, and a projection to determine the pressure that enforces incompressibility [40]. This gives us the velocity field $\boldsymbol{u}_g^{n+1/2}$, the mid-step approximation required in an RK2 method.

6. Using the mid-step fluid velocity $\boldsymbol{u}_g^{n+1/2}$ and the mid-step *data site* positions $(\boldsymbol{X}_d)_j^{n+1/2}$, update the *data sites* $(\boldsymbol{X}_d)_j^{n+1/2}$ for each j to the time level t^{n+1} using

$$(\boldsymbol{X}_d)_j^{n+1} = (\boldsymbol{X}_d)_j^n + \Delta t(\boldsymbol{U}_d)_j^{n+1/2} \equiv (\boldsymbol{X}_d)_j^n + \Delta t \sum_g \boldsymbol{u}_g^{n+1/2} \delta_h \left(\boldsymbol{x}_g - (\boldsymbol{X}_d)_j^{n+1/2} \right) h^2,$$

$$\text{(C.4)}$$

where δ_h is the same approximate δ-function we have used throughout.

7. Update the velocity $\boldsymbol{u}_g^n$ to time-level t^{n+1} using the mid-step velocity $\boldsymbol{u}_g^{n+1/2}$ and the force $\boldsymbol{f}_g^{n+1/2}$. The mid-step velocities are advected, while a Crank–Nicolson scheme is used for time-stepping the momentum equations. The pressure projection gives us the discretely incompressible velocity field $\boldsymbol{u}_g^{n+1}$.

Observe that the data sites are updated twice per time-step in the RK2 scheme, but the sample sites are only generated once. Because the data sites are typically a fraction of the number of IB points from the PL-IB method, the computational cost is significantly lower for the RBF-IB method, even factoring in the interpolation and the sample site generation.

ACKNOWLEDGEMENTS

REFERENCES

1. Peskin CS. Numerical analysis of blood flow in the heart. *Journal of Computational Physics* 1977; **25**:220–252.
2. Peskin CS, McQueen DM. Modeling prosthetic heart valves for numerical analysis of blood flow in the heart. *Journal of Computational Physics* 1980; **37**:113–132.
3. Peskin CS, McQueen DM. A three-dimensional computational method for blood flow in the heart: I. Immersed elastic fibers in a viscous incompressible fluid. *Journal of Computational Physics* 1989; **81**:372–405.
4. Fauci LJ, Fogelson AL. Truncated Newton methods and the modeling of complex immersed elastic structures. *Communications on Pure and Applied Mathematics* 1993; **66**:787–818.
5. Fauci LJ, Peskin CS. A computational model of aquatic animal locomotion. *Journal of Computational Physics* 1988; **77**:85–108.
6. Dillon R, Fauci L, Fogelson A, Gaver D. Modeling biofilm processes using the immersed boundary method. *Journal of Computational Physics* 1996; **129**:85–108.
7. Agresar G, Linderman JJ, Tryggvason G, Powell KG. An adaptive, Cartesian, front-tracking method for the motion, deformation and adhesion of circulating cells. *Journal of Computational Physics* 1998; **143**:346–380.
8. Beyer RP. A computational model of the cochlea using the immersed boundary method. *Journal of Computational Physics* 1992; **98**:145–162.
9. Miller LA, Peskin CS. When vortices stick: an aerodynamic transition in tiny insect flight. *The Journal of Experimental Biology* 2004; **207**:3073–3088.
10. Miller LA, Peskin CS. A computational fluid dynamics of 'clap and fling' in the smallest insects. *The Journal of Experimental Biology* 2005; **208**:195–212.
11. Jackson S, Nesbitt W, Kulkarni S. Signaling events underlying thrombus formation. *Journal of Thrombosis and Hameostasis* 2003; **1**:1602–1612.
12. Savage B, Saldivar E, Ruggeri ZM. Initiation of platelet adhesion by arrest onto fibrinogen or translocation on von Willebrand factor. *Cell* 1996; **84**:289–297.
13. Fogelson AL. A mathematical model and numerical method for studying platelet adhesion and aggregation during blood clotting. *Journal of Computational Physics* 1984; **1**:111–134.
14. Fogelson AL, Guy RD. Immersed-boundary-type models of intravascular platelet aggregation. *Computer Methods in Applied Mechanics and Engineering* 2008; **197**:2087–2104.
15. Fogelson A, Kuharsky A, Yu H. Computational modeling of blood clotting: coagulation and three-dimensional platelet aggregation. In *Polymer and Cell Dynamics: Multiscale Modeling and Numerical Simulations*, Alt W, Chaplain M, Griebel M, Lenz J (eds). Birkhaeuser-Verlag: Basel, 2003; 145–154.
16. Yu H. Three dimensional computational modeling and simulation of platelet aggregation on parallel computers. *Ph.D. Thesis*, University of Utah, 2000.
17. Shankar V, Wright GB, Fogelson AL, Kirby RM. A study of different modeling choices for simulating platelets within the immersed boundary method. *Applied Numerical Mathematics* 2013; **63**(0):58–77.
18. Shankar V, Wright GB, Fogelson AL, Kirby RM. A radial basis function (RBF) finite difference method for the simulation of reaction-diffusion equations on stationary platelets within the augmented forcing method. *International Journal for Numerical Methods in Fluids* 2014; **75**(1):1–22.
19. Shankar V, Wright GB, Kirby RM, Fogelson AL. A radial basis function (RBF)-finite difference (FD) method for diffusion and reaction-diffusion equations on surfaces. *Journal of Scientific Computing* 2014; **63**(3):745–768.
20. Fuselier EJ, Wright GB. A high-order kernel method for diffusion and reaction-diffusion equations on surfaces. *Journal of Scientific Computing* 2013; **56**:535–565.

21. Shankar V, Olson SD. Radial basis function (RBF)-based parametric models for closed and open curves within the method of regularized Stokeslets, 2015. DOI: 10.1002/fld.4048.
22. Peskin CS. The immersed boundary method. *Acta Numerica* 2002; **11**:479–517.
23. Yao L, Fogelson AL. Simulations of chemical transport and reaction in a suspension of cells I: an augmented forcing point method for the stationary case. *International Journal for Numerical Methods in Fluids* 2012; **69**(11): 1736–1752.
24. Fasshauer GE. *Meshfree Approximation Methods with MATLAB*, Interdisciplinary Mathematical Sciences–vol. 6. World Scientific Publishers: Singapore, 2007.
25. Wendland H. *Scattered Data Approximation*, Cambridge Monographs on Applied and Computational Mathematics, vol. 17. Cambridge University Press: Cambridge, 2005.
26. Fasshauer GE, Schumaker LL. Scattered data fitting on the sphere. In *Mathematical Methods for Curves and Surface, Vol. 2 of the Proceedings of the 4th Int. Conf. on Mathematical Methods for Curves and Surfaces, Lillehammer, Norway*, Daehlen M, Lyche T, Schumaker LL (eds). Vanderbilt University Press: Nashville Tennessee, 1998; 117–166.
27. Jetter K, Stöckler J, Ward JD. Error estimates for scattered data interpolation on spheres. *Mathematics of Computation* 1999; **68**(226):733–747.
28. Narcowich FJ, Sun X, Ward JD, Wendland H. Direct and inverse Sobolev error estimates for scattered data interpolation via spherical basis functions. *Foundations of Computational Mathematics* 2007; **7**(3):369–390.
29. Flyer N, Wright GB. Transport schemes on a sphere using radial basis functions. *Journal of Computational Physics* 2007; **226**:1059–1084.
30. Flyer N, Wright GB. A radial basis function method for the shallow water equations on a sphere. *Proceedings of the Royal Society A* 2009; **465**:1949–1976.
31. Piret C. The orthogonal gradients method: A radial basis functions method for solving partial differential equations on arbitrary surfaces. *Journal of Computational Physics* 2012; **231**(20):4662–4675.
32. Fuselier E, Wright GB. Scattered data interpolation on embedded submanifolds with restricted positive definite kernels: Sobolev error estimates. *SIAM Journal on Numerical Analysis* 2012; **50**(3):1753–1776.
33. Griffith BE. Simulating the fluid dynamics of natural and prosthetic heart valves using the immersed boundary method. *International Journal of Applied Mechanics* 2009; **1**:137–177.
34. Fauci LJ, Peskin CS. A computational model of aquatic animal locomotion. *Journal of Computational Physics* 1988; **77**(1):85–108.
35. Harlow FH, Welch JE. Numerical calculation of time-dependent viscous incompressible flow of fluid with free surface. *Physics of Fluids (1958-1988)* 1965; **8**(12):2182–2189.
36. Liu Y, Mori Y. Properties of discrete delta functions and local convergence of the immersed boundary method. *SIAM Journal on Numerical Analysis* 2012; **50**(6):2986–3015.
37. Devendran D, Peskin CS. An immersed boundary energy-based method for incompressible viscoelasticity. *Journal of Computational Physics* 2012; **231**(14):4613–4642.
38. Newren EP, Fogelson AL, Guy RD, Kirby RM. Unconditionally stable discretizations of the immersed boundary equations. *Journal of Computational Physics* 2007; **222**:702–719.
39. Fornberg B, Piret C. A stable algorithm for flat radial basis functions on a sphere. *SIAM Journal of Scientific Computing* 2007; **30**:60–80.
40. Guy RD, Fogelson AL. Stability of approximate projection methods on cell-centered grids. *Journal of Computational Physics* 2005; **203**:517–538.