

A Minimally-Resolved Immersed Boundary Model for Reaction-Diffusion Problems

Amneet Pal Singh Bhalla,^{1,2} Boyce E. Griffith,^{3,2} Neelesh A. Patankar,¹ and Aleksandar Donev^{2,*}

¹*Department of Mechanical Engineering, Northwestern University, Evanston, IL 60208*

²*Courant Institute of Mathematical Sciences,
New York University, New York, NY 10012*

³*Leon H. Charney Division of Cardiology, Department of Medicine,
New York University School of Medicine, New York, NY 10016*

We develop an immersed-boundary approach to modeling reaction-diffusion processes in dispersions of reactive spherical particles, from the diffusion-limited to the reaction-limited setting. We represent each reactive particle with a minimally-resolved “blob” using many fewer degrees of freedom per particle than standard discretization approaches. More complicated or more highly resolved particle shapes can be built out of a collection of reactive blobs. We demonstrate numerically that the blob model can provide an accurate representation at low to moderate packing densities of the reactive particles, at a cost not much larger than solving a Poisson equation in the same domain. Unlike multipole expansion methods, our method does not require analytically-computed Green’s functions, but rather, computes regularized discrete Green’s functions on the fly by using a standard grid-based discretization of the Poisson equation. This allows for great flexibility in implementing different boundary conditions, coupling to fluid flow or thermal transport, and the inclusion of other effects such as temporal evolution and even nonlinearities. We develop multigrid-based preconditioners for solving the linear systems that arise when using implicit temporal discretizations or studying steady states. In the diffusion-limited case the resulting linear system is a saddle-point problem, the efficient solution of which remains a challenge for suspensions of many particles. We validate our method by comparing to published results on reaction-diffusion in ordered and disordered suspensions of reactive spheres.

I. INTRODUCTION

Over the past decade several types of synthetic chemically-propelled micro- and nano-swimmers have been manufactured and studied in the lab [1, 2]. Modeling the behavior of suspensions of such particles requires including both the hydrodynamics of the fluid flow *and* the reaction-diffusion

*Electronic address: donev@courant.nyu.edu

transport of the chemically-reacting species around the particles. Reaction-diffusion processes are ubiquitous in chemical engineering applications and have been the subject of intense study on their own right, particularly in the diffusion-limited regime in which the rate-limiting process is the diffusion of the reactants to the reactive sites [3–6]. The fluid dynamics of active suspensions is a very active field itself [7], and methods similar to the one developed here have recently been proposed [8]. Fewer theoretical and computational studies have been done that couple the Navier-Stokes equation for the fluid velocity with a reaction-diffusion equation for the concentration of the reactants [9].

Herein, we develop a model and numerical method for such reaction-diffusion problems using an approach based on the immersed boundary method [10]. We build on recently-developed computational methods for the hydrodynamics of dilute suspensions of rigid spherical particles [11, 12]. In the *minimally-resolved* approach that we adopt, each particle is represented by a single point-like object that we will call a *blob*. The blob does not resolve the details of the particle surface, but rather, represents the effective contribution of the particle using a small number of degrees of freedom, as do particles in (stochastic) dissipative particle dynamics [13]. This is very similar to the use of multipole expansions to solve fluid-particle and reaction-diffusion problems [4, 14]. The essential difference is that in our approach we do not employ analytical Green’s functions to represent the response of the continuum model to the particles. Instead, we use a standard grid-based solver for the continuum equations to generate the required response function “on the fly”. This approach allows us to easily change the boundary conditions, including cases in which the concentration of the chemical reactants affects the fluid flow via osmo-phoretic effects [15]. Furthermore, temporal evolution (unsteadiness) and even nonlinearities can easily be added, and the resolution can be refined locally using adaptive mesh refinement techniques [16].

The method developed here can be used as the core component of a more traditional immersed-boundary method [10] for reaction-diffusion problems, in which a large number of blobs (traditionally called markers or regularized delta functions) are placed on the surface at which the reactive boundary condition is specified. In fact, by changing the number of blobs used to discretize, for example, a spherical particle, one can go from high-resolution models (e.g., hundreds of blobs per particle), to low-resolution models (e.g., 12 blobs per particle), to minimally-resolved models (a single blob per particle). In this work we focus on minimally-resolved modeling of reaction-diffusion problems in particle suspensions, extending the traditional scope and the traditional view of convergence of the immersed boundary method [11, 12]. In our approach, even though each particle is represented only by the position of its centroid, it is not appropriate to consider it a “point”

particle or an approximation (regularization) to a delta function. Rather, a blob can be thought of as a diffuse sphere that has some physical extent and interacts with the continuum in its interior.

The minimally-resolved blob approach can be used in a variety of contexts where boundary conditions are specified on the surface of particles embedded in a medium. These boundary conditions can be of the Dirichlet, Neumann, or mixed (Robin) type. Boundary conditions involving normal or tangential derivatives could be specified on the exterior boundary of the particles only (e.g., free-slip for flow problems), or as a jump condition between the interior and exterior (e.g., thermal or electrical conductivity in a dispersion of conducting particles). The specifics of how to express the surface boundary conditions to a volumetric blob condition are very problem specific and need to be carefully constructed on a case by case basis.

A. Suspensions of reactive rigid spheres

In this work we develop a blob representation for dispersions of reactive spherical particles in either two or three dimensions. The particles are taken as (possibly overlapping) rigid spheres (disks in two dimensions) occupying a region $\mathcal{B} \subset \Omega$ inside a domain Ω , and we consider the general case of a reaction that is not necessarily diffusion-limited. In the region outside of the particles $\Omega \setminus \mathcal{B}$, the diffusion of a chemical reactant is described with the diffusion equation for the concentration of the species $c(\mathbf{r}, t)$,

$$\partial_t c = \chi \nabla^2 c + s(\mathbf{r}, t) \text{ in } \Omega \setminus \mathcal{B}, \quad (1)$$

where χ is the diffusion coefficient and s is an external source/sink density. Different boundary conditions may be specified on the exterior boundary of the domain $\partial\Omega$, for example, periodic boundary conditions in the case when Ω is a torus. Let us focus on a single reactive particle. The boundary condition at the surface of the particle $\partial\mathcal{B}$ is a Robin condition relating the component of the diffusive flux along the normal vector \mathbf{n} (chosen here to point into the interior of the particle) with the concentration at the surface,

$$\chi(\mathbf{n} \cdot \nabla c) = k c \text{ on } \partial\mathcal{B}, \quad (2)$$

where k is the *surface* reaction rate [5, 17, 18]. The relative speed of reaction and diffusion is measured by the dimensionless Damköhler number $\text{Da} = (ka)/\chi$, where a is the radius of the spherical particle. Here we will follow the notation of Lu [17], and use the dimensionless constant $P = \chi/(ka) = \text{Da}^{-1}$. The reaction is diffusion-limited if $P \ll 1$, and reaction-limited if $P \gg 1$. In the diffusion-limited regime $P \rightarrow 0$, which is of most practical importance, the boundary condition for (1) becomes $c = 0$ on $\partial\mathcal{B}$.

In many practical situations one is interested in the steady-state solutions, in which case one can set $\partial_t c = 0$ in (1). Quantities such as the average or total reaction rate can be obtained by solving the steady-state reaction-diffusion problem analytically, or numerically using first-passage Monte Carlo techniques [5, 6, 17, 18] or multipole expansion techniques [4, 14]. At steady-state, the total reactive flux at the surface of a given particle, $\lambda = \chi \int_{\partial\mathcal{B}} (\mathbf{n} \cdot \nabla c) dS$, is finite and gives the effective source/sink strength of the particle. For an isolated sphere of radius a immersed in a reservoir of the reactant with concentration c_∞ in the absence of sources ($s = 0$), the analytical solution (see, for example, the Appendix of Ref. [17]) has a modified Smoluchowski form,

$$c(r) = c_\infty \left(1 - \frac{1}{1+P} \frac{a}{r} \right). \quad (3)$$

The slow r^{-1} decay of the concentration is analogous to the slow decay of hydrodynamic interactions in suspensions of spheres, and leads to nontrivial interactions between the particles in the dispersion. The total rate of consumption of the reactant is

$$\lambda = \frac{4\pi a \chi c_\infty}{(1+P)}. \quad (4)$$

Let us consider a dispersion of N spheres in a domain of volume V_Ω in the case when a uniform source density of reactant is supplied, $s(\mathbf{r}, t) = \text{const}$. The steady-state average concentration $\bar{c} = V_\Omega^{-1} \int_{\Omega \setminus \mathcal{B}} c d\mathbf{r}$ will be given by the balance condition between the total consumption rate at the surfaces of the particles and the supply in the bulk. We can define an overall reaction rate constant K via the relation [17]

$$K\bar{c} = N^{-1} \left[\chi \int_{\partial\mathcal{B}} (\mathbf{n} \cdot \nabla c) dS \right] = N^{-1} s (V_\Omega - V_\mathcal{B}),$$

where $V_\mathcal{B} = N(4\pi a^3/3)$ is the volume occupied by the particles. Denoting the volume or packing fraction $\varphi = V_\mathcal{B}/V_\Omega$, we can express K as

$$K = (4\pi a^3/3) \frac{(1-\varphi)s}{\varphi\bar{c}}.$$

In a very dilute dispersion of well-separated spheres, $\varphi \ll 1$, we can ignore particle interactions and use (3) to obtain that

$$K_0 = \lim_{\varphi \rightarrow 0} K = \frac{4\pi a \chi}{1+P}.$$

A dimensionless (normalized) reaction rate constant that measures the importance of multiparticle effects can be defined as [17]

$$\beta_P = \frac{K}{K_0} = \frac{sa^2}{\chi} \cdot \frac{(1-\varphi)}{3\varphi} \cdot \frac{(1+P)}{\bar{c}}, \quad (5)$$

and measured by calculating the average concentration \bar{c} at steady state. The dimensionless quantity β_P is independent of the particular physical units and only depends on the dimensionless number P and the microstructure of the system, i.e., on φ and the arrangement of the spheres. The value of this quantity has been obtained using first-passage techniques for a variety of regular

(lattice) arrangements of spheres and for random dispersions [3–5, 17, 18].

As a model problem on which we will validate and calibrate the blob model, let us consider the steady state solution for a finite collection of N reactive spheres of radius a in an infinite three-dimensional reservoir with concentration c_∞ far away from the particles. If one approximates each reactive sphere i with just a monopole term of strength $\lambda_i = \chi \int_{\partial\mathcal{B}_i} (\mathbf{n} \cdot \nabla c) dS$, where $\partial\mathcal{B}_i$ is the surface of the sphere, then a standard multipole expansion gives the system of equations (see Section II.B in Ref. [5])

$$\chi c_\infty = \frac{\lambda_j}{4\pi a} + \frac{1}{4\pi} \sum_{i \neq j=1}^N \frac{\lambda_i}{r_{ij}}, \quad j = 1, \dots, N,$$

where $r_{ij} = \|\mathbf{q}_i - \mathbf{q}_j\|$ is the distance between particles i and j . In matrix form we can write this as

$$\chi c_\infty \mathbf{1} = \mathcal{M}\boldsymbol{\lambda}, \quad (6)$$

where $\mathbf{1}$ is a vector of ones and the matrix \mathcal{M} has entries $\mathcal{M}_{ii} = (4\pi a)^{-1}$ and $\mathcal{M}_{ij} = (4\pi r_{ij})^{-1}$ for $i \neq j$. The slow r^{-1} decay is a characteristic signature of the inverse Poisson operator and couples the particles through their far field contributions. Note that in two dimensions the Green’s function for the Poisson equation behaves logarithmically and the boundary conditions affect the result strongly regardless of the system size. In fact, in two dimensions one cannot prescribe the concentration at infinity (this is related to the Stokes paradox for flow around a cylinder).

II. REACTIVE BLOB MODEL

In the beginning, we focus on the continuum formulation of the continuum-particle coupling. However, it is important to point out that most of the notation and conclusions can directly be adopted in the discrete formulation by simply replacing spatial integrals with sums over grid points. We will return to the spatially-discrete formulation in Section III.

A. The Blob Model of a Particle

Let us consider a single spherical reactive particle of radius a at position $\mathbf{q}(t)$ in a domain with Cartesian coordinate \mathbf{r} . The shape of the particle and its effective interaction with the surrounding medium is captured through a smooth kernel function $\delta_a(\mathbf{r})$ that integrates to unity and whose support is localized in a region of size a . In the continuum setting considered here, one may choose any one-dimensional “bell-shaped” curve $\delta_a(r)$ with half-width of order a , and define a spherically-

symmetric $\delta_a(\mathbf{r}) = \delta_a(r)$; alternatively, in d dimensions one may define a tensor-product

$$\delta_a(\mathbf{r}) = \prod_{\alpha=1}^d \delta_a(r_\alpha). \quad (7)$$

In immersed-boundary methods [10], the kernel function δ_a is considered to be an approximation of the Dirac delta function of purely numerical origin and has the tensor-product form (7). Here we choose the shape of the function based on numerical considerations, but relate its shape to the physical properties of the particle, namely, its effective reactive radius, surface and volume. This is very similar to the approach taken for particulate flows in the Force Coupling Method [19].

The reactant concentration field $c(\mathbf{r}, t)$ is extended over the whole domain, *including* the particle interior. Interaction between the field and particle is mediated by the kernel function through two crucial local operations. The *local averaging* operator $\mathbf{J}(\mathbf{q})$ averages the concentration inside the particle to estimate the local concentration

$$c_{\mathbf{q}}(t) = \int \delta_a(\mathbf{q} - \mathbf{r}) c(\mathbf{r}, t) d\mathbf{r} \equiv [\mathbf{J}(\mathbf{q})] c.$$

The reverse of local averaging is accomplished using the *local spreading* linear operator $\mathbf{S}(\mathbf{q})$ which takes a point source/sink strength $\lambda(t)$ at the location of the particle and spreads it over the extent of the kernel function to return a smooth source density field,

$$\lambda_{\mathbf{q}}(\mathbf{r}, t) = \lambda(t) \delta_a(\mathbf{q} - \mathbf{r}) \equiv [\mathbf{S}(\mathbf{q})] \lambda.$$

Note that the local spreading operator \mathbf{S} has dimensions of inverse volume. For notational simplicity we will slightly abuse notation and assume that the local spreading and interpolation operators can be applied to a scalar, a vector, or a tensor field, with the interpretation that the same local averaging or spreading operation is applied to each component independently. This sort of block-diagonal form of the spreading and interpolation operators is not strictly required for the mathematical formulation, but applies to the specific forms of the operators we use in practice [10].

The physical volume of the particle ΔV is related to the shape and width of the kernel function via $\mathbf{JS} = \Delta V^{-1} \mathbf{I}$, that is,

$$\Delta V = (\mathbf{JS} \mathbf{1})^{-1} = \left[\int \delta_a^2(\mathbf{r}) d\mathbf{r} \right]^{-1}. \quad (8)$$

Therefore, even though the particle is represented only by the position of its centroid, it is not appropriate to consider it a “point” particle. Rather, it can be thought of as a diffuse sphere that has some physical extent. For lack of better terminology, we will refer to such a diffuse particle as a “blob”.

For multi-particle problems, the positions of N blobs is described by the configuration $\mathbf{Q} = \{\mathbf{q}_1, \dots, \mathbf{q}_N\}$. We can define a composite local averaging operator $\mathcal{J}(\mathbf{Q})$ that separately averages

the concentration in the neighborhood of each particle,

$$(\mathcal{J}c)_i \equiv [\mathbf{J}(\mathbf{q}_i)] c,$$

as well as a composite local spreading operator that adds the contributions from all the sources/sinks $\boldsymbol{\lambda} = \{\lambda_1, \dots, \lambda_N\}$,

$$\mathcal{S}\boldsymbol{\lambda} = \sum_{i=1}^N [\mathcal{S}(\mathbf{q}_i)] \lambda_i.$$

It is important to note that the averaging and spreading operators are *adjoint*, $\mathcal{S} = \mathcal{J}^*$. The corresponding spatially-discrete operators (matrices) will be scaled transposes of each other. Also note that if the kernels are compactly-supported and the supports of the kernels of the different particles are not overlapping, $\mathcal{J}\mathcal{S} = \Delta V^{-1} \mathbf{I}_N$ is a multiple of the identity operator.

B. Reaction-Diffusion Equation with Blobs

Let us consider a single reactive blob. Instead of coupling the particle and the medium via surface integrals, as in (1,2), we model the reaction-diffusion system by

$$\partial_t c = \chi \nabla^2 c - \sum_i \kappa_i \left[\int \delta_a(\mathbf{q}_i - \mathbf{r}) c(\mathbf{r}, t) d\mathbf{r} \right] \delta_a(\mathbf{q}_i - \mathbf{r}) + s(\mathbf{r}, t),$$

where κ_i is an overall reaction rate near blob i . Using the local interpolation and spreading operators the reaction-diffusion equation can be succinctly written as

$$\partial_t c = \chi \nabla^2 c - \left(\sum_{i=1}^N \kappa_i \mathcal{S}_i \mathbf{J}_i \right) c + s = \chi \nabla^2 c - (\mathcal{S}\boldsymbol{\kappa}\mathcal{J}) c + s, \quad (9)$$

where $\boldsymbol{\kappa}$ is a diagonal matrix of reaction rates. Note that when particles consume reactant, $\kappa > 0$, the operator $-\mathcal{S}\boldsymbol{\kappa}\mathcal{J}$ is negative-semidefinite just like the Laplacian operator, and therefore the dynamics (9) is strictly dissipative.

The surface reaction rate k in (1,2) has physical units of m/s . The blob reaction rate $\kappa = kA \sim k\Delta V/a \sim ka^2$ has units m^3/s , where A is the surface area of the blob. In order to relate the blob model to the perfect sphere model, we take

$$\kappa = kA = 4\pi ka^2, \quad (10)$$

where a is an effective reactive radius of the blob, related to the width of the kernel function δ_a . We will justify the identification (10) in Section III D. Recall that the inverse Damköhler dimensionless number

$$P = \frac{\chi}{ka} = 4\pi \frac{\chi a}{\kappa} \quad (11)$$

determines whether the process is diffusion-limited ($P \rightarrow 0$) or reaction-limited ($P \rightarrow \infty$).

In the diffusion-limited case, the total rate of consumption of the reactant by particle i , $\lambda_i =$

$\kappa_i(\mathbf{J}_i c)$, remains constant even though $\kappa_i \rightarrow \infty$. This means that $\mathbf{J}_i c \rightarrow 0$, which corresponds to the boundary condition $c = 0$ on $\partial\mathcal{B}$ in the case of a fully-resolved reactive sphere. In the limit $P \rightarrow 0$ the reaction-diffusion equation (9) approaches the saddle-point problem

$$\begin{aligned} \partial_t c &= \chi \nabla^2 c - \mathcal{S}\boldsymbol{\lambda} + s, \\ \text{s.t. } \mathcal{J}c &= 0, \end{aligned} \tag{12}$$

where the sink strengths $\boldsymbol{\lambda} \leftarrow \kappa \mathcal{J}c$ are a Lagrange multiplier corresponding to the constraint. We remark that (9) can be viewed as a penalty method for the constrained formulation (12). Even for problems with a large but finite reaction rates, it may be preferable to use the saddle-point formulation because (9) becomes increasingly ill conditioned the limit of large κ . We also remark that effective solvers for (9) may serve as effective preconditioners for (12), although this has not yet been attempted by us.

III. SPATIO-TEMPORAL DISCRETIZATION

In this section we describe our spatial discretization of (9) and (12), and then focus on methods for efficiently solving the resulting linear systems of equations using multigrid-based preconditioning.

A. Spatial Discretization

Our spatial discretization of the equations (9) is based on standard finite-volume techniques for the Poisson equation, combined with techniques from the immersed boundary method [10]. In this work we focus on a uniform Cartesian grid of spacing h , although it is possible to extend the approach to non-uniform adaptively-refined grids [16, 20].

In the discrete setting, the various continuum operators acting on vector fields become matrices. The concentration field c is discretized as a vector \mathbf{c} of cell-centered concentrations. The Laplacian $\nabla^2 c$ is discretized as $\mathbf{L}\mathbf{c}$ using the standard second-order $(2d + 1)$ -point discrete Laplacian \mathbf{L} , along with suitable modifications near physical boundaries depending on the boundary condition imposed. Higher-order or more isotropic discrete Laplacians can be used to lessen the discretization artifacts, but at the cost of more complicated multigrid linear solvers and difficulties with conservative adaptive-mesh refinement.

Application of the local averaging operator \mathbf{J} , which is a convolution operator in the continuum

setting, becomes a discrete summation over the grid points that are near the particle,

$$\mathbf{J}\mathbf{c} \equiv \sum_{\mathbf{k} \in \text{grid}} \phi_w(\mathbf{q} - \mathbf{r}_k) c_k,$$

where \mathbf{r}_k denotes the center of the cell (control volume) with which c_k is associated. Here $\phi_w(x)$ is a *kernel* function with compact support of width wh . Note that ϕ_w is related to the kernel function δ_a that appeared in the continuum formulation, but is not in general the same function. In particular, in the discrete setting the kernel ϕ_w needs to be carefully constructed taking into account the underlying discrete grid. This grid is not rotationally invariant, and this complicates the use of spherically-symmetric kernels.

We follow the approach generally employed by immersed boundary methods [10] and do the local averaging independently along each of the d dimensions,

$$\phi_w(\mathbf{q} - \mathbf{r}_k) = \prod_{\alpha=1}^d \phi_w[q_\alpha - (r_k)_\alpha], \quad (13)$$

which improves the isotropy of the spatial discretization. Constructing kernels ϕ_w that provide both translational and rotational invariance would require not using tensor-product kernels; this difficult task has not, to our knowledge, been accomplished yet. As a matrix, the local spreading operator $\mathbf{S} = \mathbf{J}^* = (\Delta V_f)^{-1} \mathbf{J}^T$,

$$(\mathbf{S}\lambda)_k = (\Delta V_f)^{-1} \phi_w(\mathbf{q} - \mathbf{r}_k) \lambda,$$

where $\Delta V_f = h^d$ is the volume of a hydrodynamic cell. Examples of two compact-support ($w = 3$ and $w = 4$) discrete kernel functions in two dimensions are illustrated in Fig. 1 as color plots of the value $\phi_w(\mathbf{q} - \mathbf{r}_k)$ corresponding to the cell centers \mathbf{r}_k around a blob. The one dimensional kernel function used to construct the two-dimensional kernel via the tensor product (13) is shown for comparison.

The kernel function ϕ_w was constructed by Peskin [10] to yield translationally-invariant zeroth- and first-order moment conditions, along with a quadratic condition,

$$\begin{aligned} \sum_{\mathbf{k} \in \text{grid}} \phi_w(\mathbf{q} - \mathbf{r}_k) &= 1 \\ \sum_{\mathbf{k} \in \text{grid}} (\mathbf{q} - \mathbf{r}_k) \phi_w(\mathbf{q} - \mathbf{r}_k) &= 0 \\ \sum_{\mathbf{k} \in \text{grid}} \phi_w^2(\mathbf{q} - \mathbf{r}_k) &= \Delta V^{-1} = \text{const.}, \end{aligned} \quad (14)$$

independent of the position of the particle \mathbf{q} relative to the underlying (fixed) velocity grid. Ensuring these properties requires making the support of the kernel function an integer multiple w of the grid spacing h , specifically, taking $\phi_w(x) = \tilde{\phi}_w(x/h)$, where $\tilde{\phi}_w$ is compactly supported on the interval $[-w/2, w/2]$. This means that the size and shape of the particle $a \sim h$ is determined




Figure 1: Illustration of discrete kernel functions used to represent the interaction between the particles and the grid used to solve the concentration equation. The position of the blob \mathbf{q} is shown with a dot. The color of each cell with center \mathbf{r}_k corresponds to the value of the kernel $\phi_w(\mathbf{q} - \mathbf{r}_k)$ weighting that cell. The one dimensional kernel function used to construct the two-dimensional kernel via the tensor product (13) is shown as an inset. (*Left panel*) The three-point ($w = 3$) Peskin kernel φ_3 . (*Right panel*) The four-point ($w = 4$) Peskin kernel φ_4 .

by the spatial discretization of the diffusion equation and the kernel used to transfer information between the particle and the grid, and cannot be chosen arbitrarily. We will obtain the precise relationship between the reactive radius a of a single reactive blob and the grid spacing h in Section IV.

The last condition (14), was imposed by Peskin [10] as a way of approximating independence under shifts of order of the grid spacing. This property implies that the particle volume $\Delta V = (\mathbf{J}\mathbf{S})^{-1}$ will remain constant and independent of the position of the blob relative to the underlying grid. The function with minimal support that satisfies (14) is uniquely determined [10]. In some of our numerical experiments we employ this *three-point* discrete kernel function φ_3 [11, 20], illustrated in the left panel of Fig. 1. This particular choice gives $\Delta V = 2^d \Delta V_f = 2^d h^d$, where d is the dimensionality. We find, consistent with Ref. [12], better translation when using the Peskin *four-point* discrete kernel function φ_4 [10], which gives $\Delta V = (8/3)^d h^d$, and is illustrated in the right panel of Fig. 1.

B. Temporal Discretization

In this section, we describe how to integrate the spatially-discretized equations from time $n\Delta t$ to time $(n+1)\Delta t$, where Δt is the time step size. We will use a superscript to denote the time level at which a given quantity is evaluated.

As a typical temporal integrator for (9), let us consider using a backward Euler step for \mathbf{c} ,

$$\frac{\mathbf{c}^{n+1} - \mathbf{c}^n}{\Delta t} = \chi \mathbf{L} \mathbf{c}^{n+1} - \mathcal{S}^n \kappa \mathcal{J}^n \mathbf{c}^{n+1} + \mathbf{s}^n. \quad (15)$$

The implicit update (15) requires solving a linear system of the form

$$(\Delta t^{-1} \mathbf{I} - \chi \mathbf{L} + \mathcal{S} \kappa \mathcal{J}) \mathbf{c} = \mathbf{B} \mathbf{c} = \mathbf{g}, \quad (16)$$

which we discuss in the next section. The relative importance of the Laplacian term and the scaled identity term is measured by the diffusive CFL number $\vartheta = \chi \Delta t / h^2$. The Euler scheme is only first-order accurate. Second-order accuracy can be achieved by using the implicit midpoint rule for the diffusive term, however, the advantage of using the backward Euler method is that one can take very large time steps, $\Delta t \rightarrow \infty$. Namely, in the limit $\vartheta \rightarrow \infty$ the backward Euler method approaches a steady-state solver,

$$(\mathcal{S} \kappa \mathcal{J} - \chi \mathbf{L}) \mathbf{c} = \mathbf{s}. \quad (17)$$

C. Diffusion-Limited Case

In the limit $P \rightarrow 0$, we need to solve a saddle-point problem (12) for \mathbf{c} and $\boldsymbol{\lambda}$,

$$\begin{aligned} \mathbf{A} \mathbf{c} + \zeta \mathcal{S} \boldsymbol{\lambda} &= \mathbf{g} \\ \xi \mathcal{J} \mathbf{c} &= \mathbf{f}, \end{aligned} \quad (18)$$

where $\boldsymbol{\lambda} \leftarrow \zeta^{-1} \kappa \mathcal{J} \mathbf{c}$ is the unknown sink strength (rate of consumption of reactant) for each of the N blobs. Here $\mathbf{A} = \Delta t^{-1} \mathbf{I} - \chi \mathbf{L}$, and we have allowed for a nonzero \mathbf{f} since this generality is necessary when constructing preconditioners for solving the saddle-point problem. We have added two scaling constants ζ and $\xi = \zeta / \Delta V_f$ to make the system symmetric and potentially improve its numerical scaling in cases when the different physical values have vastly different units (a reasonable choice is $\zeta \sim \chi h$), as an alternative to non-dimensionalization of the equations. Notice that if we impose periodic or homogeneous Neumann boundary conditions in the steady-state case, $\mathbf{A} = -\chi \mathbf{L}$, the matrix \mathbf{A} has a null space consisting of constant vectors, and the range of \mathbf{A} is the set of vectors of mean zero. Let us first consider the case when Δt is finite and \mathbf{A} is invertible.

To construct a preconditioner for the saddle-point problem, we write down the exact solution

using a Schur complement approach. From the first equation,

$$\mathbf{c} = \mathbf{A}^{-1}(\mathbf{g} - \zeta \mathbf{S}\boldsymbol{\lambda}), \quad (19)$$

and plugging this into the constraint we get

$$\xi \zeta (\mathcal{J}\mathbf{A}^{-1}\mathbf{S}) \boldsymbol{\lambda} = \xi \mathcal{J}\mathbf{A}^{-1}\mathbf{g} - \mathbf{f} = \mathbf{h}.$$

This gives the solution for the Lagrange multipliers $\boldsymbol{\lambda}$ in terms of the inverse of the Schur complement $\mathbf{M} = \mathcal{J}\mathbf{A}^{-1}\mathbf{S}$,

$$\boldsymbol{\lambda} = (\zeta\xi)^{-1} (\mathcal{J}\mathbf{A}^{-1}\mathbf{S})^{-1}\mathbf{h} = (\zeta\xi)^{-1} \mathbf{M}^{-1}\mathbf{h}. \quad (20)$$

To understand the physical meaning of the matrix \mathbf{M} , consider a finite collection of N reactive blobs in an infinite three-dimensional reservoir with concentration c_∞ far away from the particles. Setting $\zeta = \xi = 1$ and considering steady state $\mathbf{A} = -\chi\mathbf{L}$, the excess concentration $\delta\mathbf{c} = \mathbf{c} - c_\infty$ solves the saddle-point problem

$$\begin{aligned} \chi\mathbf{L}(\delta\mathbf{c}) &= \mathbf{S}\boldsymbol{\lambda}, \\ \mathcal{J}(\delta\mathbf{c}) &= -c_\infty, \end{aligned} \quad (21)$$

the solution of which can be read from (20) to be

$$-(\mathcal{J}\mathbf{L}^{-1}\mathbf{S}) \boldsymbol{\lambda} = \mathbf{M}\boldsymbol{\lambda} = \chi c_\infty.$$

This equation can be seen as a discretization of the monopole expansion (6), with \mathbf{M} being a particular discretization of \mathcal{M} . The diagonal values will be $M_{ii} \sim h^{-1}$, to within finite-size corrections. In fact, the value of the diagonal element of \mathbf{M} can be used to define an effective reactive radius $a \sim h$ for a blob, as we will study in more detail in Section IV. For two blobs i and j separated by a distance $r_{ij} \gg h$, we have that $M_{ij} \sim r^{-1}$ just like \mathcal{M}_{ij} . When the two blobs are close to each other, the r^{-1} singularity of the Green's function for the Poisson equation is regularized by the shape of the kernel used to describe the blob. This is very similar to how the method of regularized Stokeslets [21] regularizes the singular Green's function for Stokes flow, and how the Rotne-Prager tensor or similar quasi-Gaussian regularizations are used to regularize the mobility in Brownian dynamics simulations [22]. Note that in two dimensions the decay is logarithmic in r and thus much slower than r^{-1} , making the influence of the boundaries felt throughout the whole system.

As an illustration, let us consider the problem of a single reactive sphere in an infinite three-dimensional reservoir of reactant keeping the concentration at infinity fixed at c_∞ . The steady state solution for the concentration is given in (3) with $P = 0$ and has a characteristic r^{-1} decay away from the surface of the sphere. We mimic this situation by placing a single blob near the center a large cubic box with the concentration at the boundaries fixed to c_∞ , and set $\mathbf{f} = 0$. The




Figure 2: Ratio $c(\mathbf{r})/c_\infty$ for a single blob placed at the center a large cubic box with the concentration at the boundaries fixed to c_∞ . A slice through the three-dimensional domain is shown. The four-point kernel ϕ_4 is used as the kernel. (*Left panel*) The discrete concentration field \mathbf{c} (color plot) and the center of the blob (black dot). (*Right panel*) An interpolated concentration $c(\mathbf{r})$ and the equivalent rigid reactive sphere (black disk) showing the effective reactive radius of the blob.

solution of the steady-state ($\mathbf{A} = -\chi\mathbf{L}$) diffusion-limited problem (18) is illustrated in Fig. 2. A two-dimensional slice containing the center of the blob is shown. The panel on the left shows the discrete concentration field \mathbf{c} , while the the panel on the right shows an interpolated concentration $c(\mathbf{r})$ corresponding to the discrete \mathbf{c} . It can be seen in the left panel that some of the concentrations near the center of the blob are negative, which is clearly unphysical. That this has to be so is easily seen from the fact that $\mathbf{J}\mathbf{c} = 0$ and all weights in the definition of the local averaging operator \mathbf{J} are positive. In the right panel we show the equivalent reactive rigid sphere as a black disk (the procedure used to calculate its radius a is detailed in Section IV A), and it is seen that the concentration outside of the blob kernel is positive, and in fact, we will show in Section IV A that the numerical $c(\mathbf{r})$ quantitatively matches the field around a reactive rigid sphere. This illustrates the fact that while the minimally-resolved blob approach does not resolve the details in the very vicinity of the particles, it reproduces the correct far field response.

D. Finite Reaction Case

To relate the parameters k and κ , consider the case of an isolated particle in a large periodic domain with a constant supply of reactant $s = V_\Omega^{-1}$. For a very large domain $s \rightarrow 0$ and the steady

state solution will be the same as for the case of an isolated particle in an infinite reservoir with concentration at infinity fixed at c_∞ , with the value of c_∞ determined from the balance of the consumption of reactant by the blob and the total inflow of reactant $sV_\Omega = \lambda = 1$. The condition $\lambda = 1$ together with (4) gives

$$c_\infty(P) = \frac{1}{4\pi a\chi} + \frac{P}{4\pi a\chi} = c_\infty(P=0) + \frac{P}{4\pi a\chi}.$$

For a large system the average value of concentration in the domain is the same as the far-field value c_∞ ,

$$\bar{c}_P = \bar{c}_0 + \frac{P}{4\pi\chi a}, \quad (22)$$

where \bar{c}_P denotes the average steady-state concentration at finite P , and \bar{c}_0 is the corresponding value for $P = 0$. It is this relation that can be taken as the definition of P for a blob, and (10) can be derived from this definition, as we show next.

For a blob in a periodic domain with a constant source, c_P is the solution to

$$\chi\mathbf{L}c_P - \kappa\mathbf{S}\mathbf{J}c_P + s = 0,$$

while c_0 is the solution to the saddle-point problem

$$\chi\mathbf{L}c_0 - \mathbf{S}\lambda + s = 0, \quad \mathbf{J}c_0 = 0,$$

with solution $\lambda = sV_\Omega = 1$ determined from the overall balance of reactant inflow and outflow. The difference $\tilde{c} = c_P - c_0$ is the solution to

$$\chi\mathbf{L}\tilde{c} = \mathbf{S}(\kappa\mathbf{J}\tilde{c} - \lambda).$$

For this equation to be solvable, the right hand side must have mean zero, which in this case of a single particle implies

$$\kappa\mathbf{J}\tilde{c} = \lambda, \text{ and } \mathbf{L}\tilde{c} = 0.$$

This means that \tilde{c} is a constant, $\tilde{c} = \lambda/\kappa = \kappa^{-1}$, which is consistent with (3), which gives $\tilde{c}(r) = c(r; P) - c(r; P=0) = P/(4\pi\chi a)$, if we identify $\kappa^{-1} = P/(4\pi\chi a)$. From $\bar{c}_P = \bar{c}_0 + \kappa^{-1}$ and (22) we also get (11), confirming that (10) is the correct way to relate k for a rigid sphere and κ for a blob.

E. Iterative Linear Solvers

The most challenging aspect of the numerical algorithm is the solution of linear systems such as (16), (17), and (18). We solve these linear systems using an iterative (matrix-free) Krylov solver. Because of the inherent ill-conditioning in the equations, efficient solution by iterative methods requires constructing good preconditioners (approximate solvers) for the equations. In

this section we discuss and test some ideas for constructing effective preconditioners, especially in the most difficult case of diffusion-limited steady state problems. Although this section is somewhat technical, the construction of effective solvers is necessary to treat multi-particle systems.

The solution of a Helmholtz (Poisson in the steady-state case) equation

$$(\Delta t^{-1} \mathbf{I} - \chi \mathbf{L}) \mathbf{c} = \mathbf{A} \mathbf{c} = \mathbf{g} \quad (23)$$

can be accomplished iteratively very efficiently by using standard geometric multigrid methods. We rely on the IBAMR library [16] to provide these multigrid solvers. The multigrid method consists of repeated application of multigrid V cycles, and typically each V cycle reduces the norm of the residual by a factor of about 5 – 10, depending on the boundary conditions and details of the multigrid method. An alternative method to solve the Helmholtz equation is to use an Krylov method such as preconditioned conjugate gradients (PCG). One or several cycles of multigrid can be used as a preconditioner. For the simple constant-coefficient Helmholtz or Poisson problem the overall performance of the Krylov method is similar to that of pure multigrid, however, in many cases the Krylov iteration makes the linear solver more robust. Since the computational cost is approximately proportional to the total number of multigrid cycles, we use this as a proxy for the CPU effort needed to obtain a given reduction of the residual. In all convergence plots reported here the y axis is the relative residual (on a logarithmic scale) and the x axis is the total number of Helmholtz or Poisson multigrid cycles.

1. Finite Reaction Rate

To solve (16) or (17), which are of the form $\mathbf{B} \mathbf{c} = \mathbf{g}$, we use a Krylov method preconditioned with n cycles of multigrid for the Helmholtz problem (23),

$$\mathbf{P}_n^{-1} \approx (\Delta t^{-1} \mathbf{I} - \chi \mathbf{L})^{-1}. \quad (24)$$

Note that the preconditioned linear operator $\mathbf{P}_n^{-1} \mathbf{B}$ is not symmetric. Here we employ the left-preconditioned FGMRES method as a robust Krylov solver. The FGMRES solver uses the modified Gram-Schmidt orthogonalization to avoid the potential occurrence of a degenerate Krylov basis. Note that we have tried other preconditioners, such as a multiplicative preconditioner that separately solves the pure diffusive and pure reactive sub-problems in sequence, but have found no substantial improvement over the simple preconditioner (24).

We study the performance of this preconditioner on the steady state equation (17) for a periodic cubic array of reactive blobs. The periodic cubic domain is discretized into L^3 cells, and blobs are placed 4 cells apart along each of the coordinate axes, making the total number of blobs $(L/4)^3$. For




Figure 3: Convergence history of the iterative linear solver for (17) for a cubic array of $(L/4)^3$ blobs in a periodic domain of L^3 cells. A single V cycle of multigrid for the Poisson equation is used as a preconditioner, and a restart frequency of 30 is used in the FGMRES algorithm. The relative norm of the residual versus the total number of multigrid cycles (a proxy for the computing effort) is shown.

these tests we employ the 4-point kernel, so that the kernels of the blobs are just touching but not overlapping. Let us define the dimensionless number $\tilde{P} = (4\pi) \chi h / \kappa \sim P$ assuming that $a \approx h$ in (11). In Fig. 3 we report the performance of the multigrid preconditioner (24) for $\tilde{P} = (4\pi) \cdot 10^{-4}$, which is in the diffusion-limited regime, for several system sizes L ranging from 32 to 128 grid cells. We show the relative residual of the linear solver as a function of the number of multigrid V cycles, which is equivalent to the number of FGMRES iterations since we use a single V cycle as a preconditioner, $n = 1$. For comparison, in Fig. 3 we show the performance of the iterative solver for the case of $L = 128$ and a reaction-limited regime $\kappa = 0$, $\tilde{P} \rightarrow \infty$. In this case the linear system (17) is a simple Poisson problem and multigrid is known to be an excellent iterative solver. We find similarly good performance for $\tilde{P} = 1$, as also shown in the figure.

Figure 3 shows that the multigrid preconditioner is only effective for smaller system sizes or for non-diffusion-limited systems, $\tilde{P} \gtrsim 1$. For $\tilde{P} \ll 1$ the solver shows slow convergence for $L = 128$, and we have not yet found any effective simple preconditioner. Recently, specialized multigrid solvers have been proposed for solving (17) [23, 24]. In these methods, a matrix representation of the reaction operator $\mathcal{S}\kappa\mathcal{J}$ is formed, and coarsened versions of this operator are constructed using standard geometric multigrid techniques. These methods show excellent promise and in future work

we will explore their application to reaction-diffusion problems. Because these solution techniques [23, 24] consider the case of finite κ , it is not clear whether they will aid in the solution of the diffusion-limited case $\kappa \rightarrow \infty$, in which case it is appropriate to use a saddle-point formulation.

2. Diffusion-Limited Case

In the diffusion-limited case $P = 0$, we need to solve the saddle-point system (18). Formally, the solution can be found from the inverse of the Schur complement $\mathbf{M} = \mathcal{J}\mathbf{A}^{-1}\mathbf{S}$ via (19,20). In the steady-state case, \mathbf{M} is a discretization and regularization of the monopole matrix \mathcal{M} .

In multipole expansion methods, the $N \times N$ dense matrix \mathcal{M} would be constructed (though not necessarily assembled), and then a linear system such as (6) solved. Sophisticated fast-multipole methods can be used to obtain the action of \mathcal{M} with cost $O(N \log N)$ or even $O(N)$ [4, 14, 22]. With different boundary conditions, the Green's function for the Poisson problem would have to be obtained analytically or by using boundary-integral methods [25]. In the blob approach, the Green's function is approximated with \mathbf{L}^{-1} by solving a discretization of the Poisson equation on a grid. This is simpler to implement and allows for much larger flexibility in handling boundary conditions, however, it requires the development of an approximation of the action of \mathbf{M}^{-1} in order to construct an effective preconditioner. In recent years fast direct solvers, related to fast-multipole methods, have been developed to efficiently calculate the action of \mathcal{M}^{-1} , however, these methods are quite complex and still an active area of research. Once developed such methods could be used to approximate the action of \mathbf{M}^{-1} .

Here we take a different approach. We use a Krylov solver to solve the block saddle-point system (18), and use an approximation to the inverse of the Schur complement \mathbf{M}^{-1} to construct a preconditioner for this Krylov solver. The simplest possible approximation to the Schur complement matrix is to use a diagonal preconditioner based on a single-blob approximation. Denote the *scalar* value

$$\gamma = -(\mathbf{J}\mathbf{L}^{-1}\mathbf{S}) \mathbf{1},$$

which depends on the dimensionality and the system size, and, to some extent, on the position of the blob relative to the underlying grid. In three dimensions $\gamma \approx (4\pi a)^{-1}$ for large systems, where $a \sim h$ is an effective reactive radius of a blob that we calculate in Section IV. A diagonal approximation to the Schur complement,

$$\mathbf{M}^{-1} \approx \chi \gamma^{-1} \mathbf{I}, \tag{25}$$

can be used together with (20) to approximate

$$\boldsymbol{\lambda} \approx \chi (\zeta \xi \gamma)^{-1} \mathbf{h}.$$

We can then obtain an approximation to \mathbf{c} from (19) using an inexact computation of \mathbf{A}^{-1} . Specifically, we approximate $\mathbf{A}^{-1} \approx \tilde{\mathbf{A}}_n^{-1}$ by using n cycles of geometric multigrid, both when computing \mathbf{h} , and when computing \mathbf{c} . We will refer to the preconditioner obtained in this manner as the *diagonal preconditioner* even though the preconditioner matrix itself is not diagonal. Note that each Krylov iteration of the diagonal preconditioner requires $2n$ multigrid cycles. A more detailed description of the preconditioner steps is given in Appendix A.

Another approach is to approximate $\mathbf{M}^{-1}\mathbf{h}$ with an inexact Krylov solver for solving $\mathbf{M}\boldsymbol{\lambda} = \mathbf{h}$. This *inner* Krylov solver estimates $\boldsymbol{\lambda}$, which is in turn used to estimate \mathbf{c} from (19) using an inexact computation of \mathbf{A}^{-1} . These approximate values are used to construct a preconditioner for the outer Krylov solver, which solves the saddle-point system (18). This *approximate Schur complement preconditioner* has two integer parameters, n and m . The number of multigrid cycles in the approximate Poisson solves in the outer Krylov iteration is $2n$, and the number of cycles in the inner Krylov solver is m , making the total number of multigrid cycles per outer Krylov iteration $m + 2n$. A more detailed description of the preconditioner steps is given in Appendix A.

We study the performance of the diagonal and approximate Schur complement preconditioners on the saddle-point problem (18) for a periodic cubic array of reactive blobs. The periodic cubic domain is discretized into L^3 cells, and blobs are placed 4 cells apart along each of the coordinate axes, making the total number of blobs $(L/4)^3$. For these tests we employ the 4-point kernel, so that the kernels of the different blobs are just touching but not overlapping. We have performed detailed investigations of how the number of multigrid cycles employed in the approximate Poisson solves m and n affect the performance of the preconditioners; here we summarize our main observations.

For the first test we take $\mathbf{A} = L^{-2}\mathbf{I} - \mathbf{L}$, which corresponds to a very large diffusive CFL number, approaching the steady-state problem. Note that adding a small multiple of the identity to the Laplacian makes \mathbf{A} invertible even with periodic boundary conditions. For the diagonal preconditioner we have found that a single multigrid cycle is optimal, $n = 1$ (data not shown). We have also observed that a large restart frequency is required in the FGMRES iteration, making the memory requirements substantial. The convergence history of the FGMRES solver is shown in Fig. 4 as a function of the system size, for both two and three dimensions.

For the approximate Schur complement preconditioner, we have found that a large m is beneficial in speeding the outer Krylov solver convergence rate, but does not help decrease the overall




Figure 4: Performance of the diagonal preconditioner with $n = 1$ for the saddle-point problem (18) with $\mathbf{A} = L^{-2}\mathbf{I} - \mathbf{L}$, as a function of the number of cells L along each dimension, in both three dimensions (left panel) and two dimensions (right panel). The restart frequency is set to set to 200 in the FGMRES solver.

computing time (i.e., total number of multigrid cycles). This is illustrated in the top left panel of Fig. 5. Increasing n was also not found to be beneficial (data not shown). Empirically we observe that a reasonable choice is to use $m = 5$ and $n = 1$. In the top right panel of Fig. 5 we show the convergence history of the FGMRES solver in three dimensions. Similar behavior is seen as for the diagonal preconditioner, but the total computational cost is somewhat increased.

In the bottom two panels of Fig. 5 we show the convergence history of the FGMRES solver we study the steady-state case $\mathbf{A} = -\mathbf{L}$ in both two and three dimensions. We utilize the approximate Schur complement modified to take into account the null space of the discrete Laplacian, as described in Appendix A. We find similar behavior as for the case of a non-singular \mathbf{A} , demonstrating that the null space has been handled properly in the preconditioner.

Figs. 4 and 5 show that neither the diagonal or the approximate Schur complement preconditioners give a solver that is robust with respect to the system size. The number of multigrid sweeps is seen to increase strongly as the size of the system is increased. This is unlike multigrid solvers for the Poisson equation, where only a mild (at most logarithmic) dependence on system size is observed. Nevertheless, the convergence is seen to be uniform and both preconditioners provide a viable iterative solver, especially for small numbers of blobs. In particular, if the number of blobs is kept small, we find that the total number of multigrid cycles does not increase as the system size increases, as illustrated in Fig. 6. The approximate Schur complement preconditioner is found to be more robust to restarts in the FGMRES algorithm, and thus requires substantially less memory than the diagonal preconditioner.




Figure 5: Performance of the approximate Schur complement preconditioner for the saddle-point problem (18). For the top row $\mathbf{A} = L^{-2}\mathbf{I} - \mathbf{L}$ and for the bottom row $\mathbf{A} = -\mathbf{L}$. The restart frequency is set to set to 30 in the FGMRES solver. (*Top Left*) Investigation of the optimal m for $n = 1$ for a three-dimensional system of 16^3 cells. (*Top Right*) System size dependence for $m = 5$ and $n = 1$. (*Bottom left*) Performance of the special handling of the null space of \mathbf{A} in three dimensions for $m = 5$ and $n = 1$. (*Bottom right*) Same as bottom left panel but in two dimensions.

IV. RESULTS

In this section we apply the reactive blob method to model reaction-diffusion problems for ordered and random dispersions of reactive spheres. We will compare the results obtained from the blob model to those obtained by other methods [3–5, 17, 18] to determine the fidelity yielded by our blob model in approximating rigid reactive spheres. A similar study was performed in the context of fluid flow problems in Refs. [12, 26] and it was found that a blob can provide a surprisingly good approximation to a rigid sphere.




Figure 6: Performance of the diagonal preconditioner with $n = 1$ (left panel), and the approximate Schur complement preconditioner with $n = 1$, $m = 5$ (right panel) for the saddle-point problem (18) with $\mathbf{A} = L^{-2}\mathbf{I} - \mathbf{L}$. Here the number of blobs, arranged on a cubic lattice, is kept fixed at $64 = 4^3$ as the number of cells L along each dimension is varied.

A. Reaction radius of a blob

We expect a blob to behave similarly to a reactive sphere of radius $a \sim h$. The constant of proportionality gives the effective reactive radius of the blob, and depends on the discrete Laplacian \mathbf{L} used to discretize the Poisson equation, on the kernel function φ used to implement the discrete local averaging and spreading operators (see Section III A), and, to a minor extent, on the position of the blob relative to the grid of cells. Here we calculate the effective reactive radius for the 3-point and 4-point kernel functions [10] and the standard 7-point discrete Laplacian in three dimensions. It is also possible to define an effective reactive radius in two dimensions. However, in two dimensions the behavior of a collection of reactive particles is very system size and boundary condition-dependent, and therefore we focus here on the more practically-relevant case of three dimensional space.

We consider a single blob in a periodic domain Ω of volume $V_\Omega = L^3 h^3$, which is equivalent to considering an infinite cubic array of reactive blobs. We focus on the diffusion-limited case $P \rightarrow 0$, and study the steady state in the presence of a uniform source of reactant of strength s throughout the domain. We solve the saddle-point problem (18) for $\xi = \zeta = 1$, $\mathbf{g} = s$, and $\mathbf{f} = 0$ using the approximate Schur complement preconditioner described in Appendix A, with $m = 5$ and $n = 1$.

For a single reactive sphere, the normalized reaction rate (5) is given by

$$\beta_0 = \frac{sa^2}{\chi} \cdot \frac{(1 - \varphi)}{3\varphi} \cdot \frac{1}{\bar{c}}, \quad (26)$$




Figure 7: Reaction radius of a blob for Peskin’s 4-point and 3-point kernels. (*Left*) The volume-fraction-dependent reaction radius as a function of inverse system size, together with fits of the form (29) with a as an unknown parameter. (*Right*) The translational invariance of the reactive blob radius (extrapolated to infinite system size using (29)), as the blob is displaced by a distance Δ relative to the center of a grid cell along several different directions.

where $\varphi = (4\pi a^3/3)/V_\Omega$ is the volume fraction (packing density) of the cubic lattice. As $\varphi \rightarrow 0$, $\beta_0 \rightarrow 1$; therefore, by measuring the average concentration \bar{c} in the domain (recall that for a blob the domain includes all cells, including those overlapping the blob) for very small volume fractions, we can obtain the reactive radius a . If we take $s = 1/V_\Omega$ we see that for $\varphi \ll 1$,

$$\beta_0 \approx (4\pi a \chi \bar{c})^{-1} \approx 1. \quad (27)$$

We define a volume fraction-dependent effective reactive radius $a_\varphi = [4\pi \chi \bar{c}(\varphi)]^{-1}$, and measure it numerically for a blob for different system sizes. In the left panel of Fig. 7 we show a_φ/h as a function of $L^{-1} \sim \varphi^{1/3}$. Theoretical calculations based on multipole expansions [27] give the asymptotic expansion

$$\beta_0 = 1 + 1.76\varphi^{1/3} + (1.76)^2 \varphi^{2/3} + \text{h.o.t.}, \quad (28)$$

which gives the corresponding expansion of a_φ , expressed in terms of system size as

$$a_L = a \left[1 + 2.84 \frac{a}{Lh} + \left(2.84 \frac{a}{Lh} \right)^2 + \text{h.o.t.} \right]. \quad (29)$$

By fitting the numerical data to this series expansion at small φ we can obtain a numerical estimate for a . It is important to note that the precise value of a depends on the position of the center of the blob relative to the underlying grid. For a blob at the node (corner) of the grid, we numerically estimate $a \approx 1.27h$ for the 4-point kernel and $a \approx 0.885h$ for the 3-point kernel.

Next we explore how translationally-invariant the effective reactive radius a is. We fix the system size at $L = 128$ and then move the blob relative to the center of a grid cell in several test




Figure 8: Decay of the concentration around a blob in a cubic domain of 100^3 grid cells with Dirichlet boundary conditions, compared to the theoretical prediction (30). The four-point kernel ϕ_4 is used so the reactive radius is taken to be $a \approx 1.255h$.

directions, and estimate the infinite system size a from (29) as $a \approx a_L - 2.84a_L^2 / (Lh)$. In the right panel of Fig. 7 we show a as a function of the distance Δ of the blob to the nearest cell center. We see that a varies only by about 1.5% for the 4-point kernel, and by about 5% for the 3-point kernel, consistent with similar results for flow problems [11, 12]. The 4-point function is seen to provide much improved translational invariance, and is therefore preferred except in cases where computational cost is of primary concern. In future work we will study other choices for the kernel functions.

The effective reaction radius is a property of the discretized blob (specifically, the combination of the solver for the concentration equation and the discrete kernel function), and does not depend on the boundary conditions. To see this, let us consider the example of a single reactive blob in a large cubic box of size L with Dirichlet boundary conditions $c = c_\infty$ at the boundaries of the domain, as first illustrated in Fig. 2. In an infinite domain the steady-state concentration would be given by (3). Here we focus on the diffusion-limited case $P = 0$. We can approximate the finite-size effects of the boundaries by solving the spherically-symmetric problem $\nabla^2 c = 0$ subject to $c(r = a) = 0$ and $c(r = L/2) = c_\infty$, to obtain

$$\frac{c(r)}{c_\infty} \approx \frac{L}{L - 2a} \left(1 - \frac{a}{r}\right). \quad (30)$$

In Fig. 8 we compare the numerical solution for the steady-state concentration around the blob

with the theoretical result, with the value of the reactive radius a taken to the average in Fig. 7. Excellent agreement is observed and the r^{-1} decay of the concentration around the blob is clearly illustrated.

B. Periodic Dispersions

We now investigate how well a periodic dispersion of reactive blobs approximates the behavior of a periodic dispersion of reactive spheres of radius a . We consider a single blob in a periodic domain Ω of volume $V_\Omega = L^3 h^3$, which is equivalent to considering an infinite cubic array of reactive blobs.

In the first test we focus on the diffusion-limited case $P \rightarrow 0$ and solve the steady-state (18) system in the presence of a uniform source of reactant of strength s throughout the domain. We obtain the normalized effective reaction rate from the average concentration in the domain using (26). The numerical results are shown in Fig. 9, for both the 3-point and 4-point kernels. Based on the theoretical prediction (28), we fit the numerical data of Lu [17] for a cubic lattice of spheres at packing fraction $\varphi = 0.1$ and $\varphi = 0.2$ to

$$\beta_0 = 1 + 1.76\varphi^{1/3} + (1.76)^2 \varphi^{2/3} + b\varphi + c\varphi^{4/3}, \quad (31)$$

with fitting (interpolation) parameters which we estimated to be $b \approx -0.92$ and $c \approx 17.4$. We take this fit as a good approximation to the true answer for a dispersion of spheres, and compare our numerical data for a dispersion of blobs to this approximation in Fig. 9. An excellent match is observed over this range of $0 < \varphi < 0.15$.

Note that achieving higher packing densities with blobs is not possible for a cubic arrangement without overlapping the kernels of the blobs. This is because a cubic arrangement is a very low-density packing of spheres. For cubic arrangements, the blob model gives unphysical results, including negative β_0 , at higher packing densities. This is because inside the support of the kernel the concentration is negative (in order to satisfy the constraint $\mathcal{J}\mathbf{c} = 0$, as illustrated in Fig. 2), and averaging the concentration over the whole domain (including the blob interiors), can produce a negative number. It is not surprising that the minimally-resolved blob model cannot provide a good approximation for densely-packed spheres. How good of an approximation the blob model provides depends, in fact, not just on the packing density but also on the type of arrangement. In the next section we study random dispersions and find that the blob model gives reasonably accurate results to rather high packing densities.




Figure 9: The effective reaction rate β_0 for a cubic array of blobs at several volume fractions, for the 3-point (left panel) and the 4-point (right panel) kernels. A theoretical prediction for a cubic array of spheres based on (31) is shown for comparison.

C. Random Dispersions

We now turn to random dispersions of blobs and compare to the first-passage results of Lee *et al.* [3] for random dispersions of non-overlapping spheres. We generate configurations of non-overlapping hard spheres at several volume fractions using the packing algorithm and code described in Refs. [28, 29]. The computational grid is kept at $L^3 = 128^3$ cells, and the number of particles is chosen to give (approximately) a desired packing fraction in the range $0.1 \leq \varphi \leq 0.4$ (the densest dispersion had on the order of 10^5 particles). Blobs are placed at the center of each hard sphere, and the steady-state diffusion-limited equation (18) is solved numerically. In the left panel of Fig. 10 we compare the data for blobs to the data for rigid spheres [3]. An excellent agreement is observed even for packing fractions as high as $\varphi = 0.4$, surprisingly close to the jamming density $\varphi \approx 0.64$. This unexpected accuracy is perhaps owing to cancellation of errors due to the randomness. Note that at higher packing densities a large fraction of the blob kernels are overlapping even though the equivalent rigid spheres would not be.

D. Finite reaction rate

The results of the previous section were obtained for diffusion-limited case $P \rightarrow 0$. Here we consider the same setup of a single blob in a periodic domain, but consider the steady state for a finite reaction rate k . Using (10), the dimensionless number $P = \chi / (ka)$ for a sphere of radius a , becomes $P = 4\pi\chi a / \kappa$ for a blob, where a is the reactive radius of the blob.




Figure 10: (Left) The effective reaction rate β_0 for a random arrangement of non-overlapping spheres at several volume fractions, for the 4-point kernel. The first-passage results of Lee *et al.* [3] are shown for comparison. (Right) The normalized inverse reaction rate Ω as a function of the dimensionless number P for a single blob in a large cubic periodic box, compared with the theoretical prediction $\Omega = \beta_0^{-1} + (1 - \varphi)P$.

Lu [30] has shown that at small volume fractions the effective reaction rate for a finite P can be related to the reaction rate for $P = 0$ via a simple approximation,

$$\beta_P = \frac{1 + P}{\Omega} \approx \frac{(1 + P)}{\beta_0^{-1} + (1 - \varphi)P}, \quad (32)$$

where φ is the volume fraction and β_0 is the normalized reaction rate for the same configuration of the spheres in the diffusion-limited case. Here

$$\Omega = \left[\frac{sa^2}{\chi} \cdot \frac{(1 - \varphi)}{3\varphi} \cdot \frac{1}{c} \right]^{-1} \approx \beta_0^{-1} + (1 - \varphi)P \quad (33)$$

can directly be measured from the steady state average concentration in the domain. Equations (32,28) give us a theoretical prediction for β_P for a cubic array of spheres at low packing density.

We numerically solve for the steady-state concentration around a single reactive blob at several finite reaction rates, using the 4-point kernel. We take a box of size $L = 100$ grid cells, which makes the finite-size effects rather small. The linear relationship (33) between Ω and P is shown in the right panel of Fig. 10 and seen to be in excellent agreement with the numerical data over the whole range of P values, confirming that the blob model captures the effect of finite reaction rate accurately.

In Fig. 11 we examine the particular case of finite reaction rate, $P = 0.1$ for a cubic lattice of reactive blobs. We compare the blob results to the theoretical prediction (32,31) for several packing fractions. A very good agreement is observed, although with some clear deviations from the theory and the numerical data of Lu [17] at the larger packing fractions. These results confirm that the blob model provides an effective approximation at low and moderate packing densities.




Figure 11: The effective reaction rate $\beta_{0.1}$ for a cubic array of spheres at several volume fractions, for the 3-point (left panel) and the 4-point (right panel) kernels. A theoretical prediction based on (32,31) is shown for comparison.

V. CONCLUSIONS

We developed a diffuse particle approach to modeling reaction-diffusion processes in particle dispersions that is based on the immersed boundary method [10] and can be used from the diffusion-limited to the reaction-limited setting. Building on positive experience with recently-developed methods for low-density colloidal suspensions [11, 12], we represented each reactive particle with a diffuse object termed a blob. In this minimally-resolved approach, each three-dimensional blob interacts with less than ~ 50 grid points ($3^3 = 27$ for the three-point kernel and $4^3 = 64$ for the four-point kernel), unlike more traditional boundary-integral, traditional immersed-boundary methods, or other discretizations in which the particle surface is explicitly resolved.

A reactive blob, just like a reactive sphere, is characterized only by a reactive radius. This reactive radius is proportional to the width of the kernel function used in the immersed boundary handling of the interaction between the blob and the computational grid used to solve Poisson's equation for the reactant concentration. For the standard kernel functions employed here the reactive radius is proportional to and, in fact, close to the grid spacing. In the future, we will explore the construction of kernels with variable width, as necessary to represent polydisperse dispersions of particles. Adaptive mesh refinement can be used in cases of very large polydispersity, or in cases when the density of blobs is strongly heterogeneous.

By comparing to accurate results in the literature, obtained using computationally expensive Monte Carlo techniques, we demonstrated numerically that the blob model can provide an accurate

representation at low to moderate packing densities of the reactive spheres, at a cost similar to solving several Poisson equations in the same domain. The blob model was shown to be effective at both finite reaction rates and the diffusion-limited case of infinite reaction rate.

At the level of the formulation, the blob method is very similar to the method of multipole expansions [4], truncated at the leading-order (monopole) level and regularized to avoid singularities. An essential difference is that our method does not require analytically-computed Green’s functions, but rather effectively computes regularized discrete Green’s functions “on the fly” by using a standard grid-based discretization of the Poisson equation. This allows for great flexibility in implementing different boundary conditions, coupling to fluid flow (including immersed-boundary methods for flow problems), and the inclusion of other effects such as temporal evolution, nonlinearities, and other terms not easily handled in a Green’s function-based approach.

More complicated particle shapes can be built out of a collection of reactive spheres, as done, for example, in Refs. [26, 31] for flow problems around rigid bodies. In fact, a more accurate representation of a spherical particle can be achieved by representing it with a collection of blobs (regularized monopoles). An alternative approach is to include dipole terms in the blob description by adding additional degrees of freedom to each blob [4]. Such an approach has been proposed by Maxey and collaborators in the context of fluid flow problems [32], and can be extended to reaction-diffusion problems in the diffusion-limited regime. At the dipole level, in addition to the unknown sink strength λ , one would associate with each particle an unknown dipole strength (a vector) $\mathbf{\Lambda}$, which corresponds to the continuum dipole moment [5]

$$\mathbf{\Lambda} = \chi \int_{\partial\mathcal{B}} \mathbf{n} \cdot (\mathbf{r} \otimes \nabla c - c\mathbf{I}) dS.$$

In the blob formulation $\mathbf{\Lambda}$ would be a Lagrange multiplier associated with a “rigidity” constraint that imposes that the concentration be constant (alternatively, that the concentration gradient be zero) in the interior of the blob [32]. Solving the extended linear system for \mathbf{c} , λ and $\mathbf{\Lambda}$ would require the development of novel linear-algebra techniques.

At present, linear algebra is a remaining bottleneck in scaling blob calculations to very large number of particles in the diffusion-limited regime. In this work we proposed and studied simple preconditioners based on using standard geometric multigrid techniques for Poisson problems as a preconditioner to a Krylov iterative solver for the saddle-point problem (18). This was found to work well for small numbers of blobs, but the conditioning number increases as the number of blobs increases, a signature of the long-ranged interactions between the blobs. An alternative approach, which we hope to explore in the future, is to use recently-developed specialized geometric multigrid

solvers for implicit immersed-boundary methods [23, 24]. In the diffusion-limited regime, it may be possible to use these multigrid solvers with a suitably chosen finite value of κ as a preconditioner for the saddle-point system. It remains to be determined how effective such preconditioners are in solving linear systems arising in the commonly-occurring case of particles interacting through a continuum via long-ranged power-law interactions.

Acknowledgments

A. Donev was supported in part by the Office of Science of the U.S. Department of Energy through Early Career award DE-SC0008271 and by the Air Force Office of Scientific Research under grant number FA9550-12-1-0356. B. Griffith acknowledges research support from the National Science Foundation under awards OCI 1047734 and DMS 1016554. A. Bhalla and N. Patankar acknowledge research support from the National Science Foundation (NSF awards CBET-0828749, CBET-1066575, and CMMI-0941674 awarded to N. Patankar). Computational resources were provided by Northwestern University High Performance Computing System – Quest.

Appendix A: PRECONDITIONING THE SADDLE-POINT SOLVER

In this Appendix we give a detailed description of the algorithm used to precondition the saddle-point system (18). The preconditioner can be thought of as an approximate solver for that linear system.

For a periodic system at steady state, $\mathbf{A} = -\chi\mathbf{L}$ is not invertible because all vectors in the range of \mathbf{A} have zero average value, $\langle \mathbf{A}\mathbf{c} \rangle = 0$ for any \mathbf{c} , where $\langle \cdot \rangle$ denotes a spatial average over the domain, $\langle \mathbf{c} \rangle = N_c^{-1} \sum_{k=1}^{N_c} c_i$, and N_c is the number of cells in the grid. In this case we need to modify the Schur complement approximation. From the fact that $\mathbf{A}\mathbf{c}$ has zero mean, we know that

$$\zeta \langle \mathbf{S}\boldsymbol{\lambda} \rangle = \langle \mathbf{g} \rangle,$$

Spreading preserves the sum of values because of the normalization condition on the kernel function,

$$\langle \mathbf{S}\boldsymbol{\lambda} \rangle = N_c^{-1} \Delta V_f^{-1} \sum_{i=1}^N \lambda_i = N V_\Omega^{-1} \langle \boldsymbol{\lambda} \rangle,$$

where $\langle \boldsymbol{\lambda} \rangle$ denotes an average over all blobs, N is the number of blobs, ΔV_f is the volume of a grid cell, and V_Ω is the volume of the domain. This means that an additional condition for solvability is

$$\langle \boldsymbol{\lambda} \rangle = V_\Omega \zeta^{-1} N^{-1} \langle \mathbf{g} \rangle,$$

which is to be added to the saddle-point problem (18) as an additional constraint. This augmented saddle-point problem has a unique solution and can be solved using the FGMRES algorithm.

Let us set $\alpha = 1$ when \mathbf{A} has a null-space, and $\alpha = 0$ when \mathbf{A} is invertible. When $\alpha = 1$, let us define the restricted inverse \mathbf{A}^{-1} to only act on vectors of mean value zero, and to return a vector of mean zero. The approximate Schur complement preconditioner and diagonal preconditioners can be implemented as follows:

1. Compute a concentration estimate

$$\tilde{\mathbf{c}}^* = \tilde{\mathbf{A}}_n^{-1} \tilde{\mathbf{g}} = \tilde{\mathbf{A}}_n^{-1} (\mathbf{g} - \alpha V_\Omega N^{-1} \mathcal{S}(\mathbf{g})),$$

where the notation $\mathcal{S}q$ denotes \mathcal{S} applied to a vector of length N with all entries equal to q . Here $\tilde{\mathbf{A}}_n^{-1}$ is an approximation of \mathbf{A}^{-1} obtained by using n cycles of geometric multigrid, starting from a zero initial guess.

2. Solve the Schur complement system

$$\left(\mathcal{J} \tilde{\mathbf{A}}^{-1} \mathcal{S} \right) \tilde{\boldsymbol{\lambda}} = (\zeta \xi)^{-1} \tilde{\mathbf{h}} = (\zeta \xi)^{-1} (\xi \mathcal{J} \tilde{\mathbf{c}}^* - \mathbf{f}), \quad (\text{A1})$$

approximately. If $\alpha = 1$, ensure that $\langle \tilde{\boldsymbol{\lambda}} \rangle = 0$ in the end by subtracting the mean.

- (a) For the diagonal preconditioner, the Lagrange multipliers would be estimated as

$$\tilde{\boldsymbol{\lambda}} \approx \chi (\zeta \xi \gamma)^{-1} \tilde{\mathbf{h}}.$$

- (b) For the approximate Schur complement preconditioner, use m iterations of a Krylov method to solve (A1), approximating $\tilde{\mathbf{A}}^{-1} = \tilde{\mathbf{A}}_1^{-1}$ by using a single cycle of geometric multigrid.

3. Correct the concentration

$$\tilde{\mathbf{c}} = \tilde{\mathbf{A}}_n^{-1} (\mathbf{g} - \zeta \mathcal{S} \tilde{\boldsymbol{\lambda}}),$$

using the previous estimate $\tilde{\mathbf{c}}^*$ as an initial guess for the multigrid solver.

4. If $\alpha = 1$, determine the average concentration $\bar{c} = \langle c \rangle$ from the condition $\langle \xi \mathcal{J} \tilde{\mathbf{c}} \rangle = \langle \mathbf{f} - \xi \bar{c} \rangle$,

$$\bar{c} = \langle \xi^{-1} \mathbf{f} - \mathcal{J} \tilde{\mathbf{c}} \rangle.$$

5. Return the approximation to the solution of the saddle-point problem,

$$\boldsymbol{\lambda} = \tilde{\boldsymbol{\lambda}} + \alpha V_\Omega \zeta^{-1} N^{-1} \langle \mathbf{g} \rangle$$

$$\mathbf{c} = \tilde{\mathbf{c}} + \alpha \bar{c}.$$

-
- [1] Raymond Kapral. Perspective: Nanomotors without moving parts that propel themselves in solution. *J. Chem. Phys.*, 138:020901, 2013.
- [2] Jeremie Palacci, Stefano Sacanna, Asher Preska Steinberg, David J Pine, and Paul M Chaikin. Living crystals of light-activated colloidal surfers. *Science*, 339(6122):936–940, 2013.
- [3] Sang Bub Lee, In Chan Kim, C. A. Miller, and S. Torquato. Random-walk simulation of diffusion-controlled processes among static traps. *Phys. Rev. B*, 39:11833–11839, Jun 1989.
- [4] R.T. Bonnecaze and John F. Brady. Rate of diffusion-limited reactions in dispersions of spherical traps via multipole scattering. *J. Chem. Phys.*, 94(1):537–540, 1991.
- [5] Heng-Kwong Tsao, Shin-Yuan Lu, and Chin-Yao Tseng. Rate of diffusion-limited reactions in a cluster of spherical sinks. *J. Chem. Phys.*, 115:3827, 2001.
- [6] A. Donev, V. V. Bulatov, T. Oettel, G. H. Gilmer, B. Sadigh, and M. H. Kalos. A First-Passage Kinetic Monte Carlo Algorithm for Complex Diffusion-Reaction Systems. *J. Comp. Phys.*, 229(9):3214–3236, 2010.
- [7] Donald L Koch and Ganesh Subramanian. Collective hydrodynamics of swimming microorganisms: Living fluids. *Annual Review of Fluid Mechanics*, 43:637–659, 2011.
- [8] Enkeleida Lushi and Charles S Peskin. Modeling and simulation of active suspensions containing large numbers of interacting micro-swimmers. *Computers & Structures*, 122(0):239 – 248, 2013.
- [9] Magdalena M Musielak, Lee Karp-Boss, Peter A Jumars, and Lisa J Fauci. Nutrient transport and acquisition by diatom chains in a moving fluid. *Journal of Fluid Mechanics*, 638:401, 2009.
- [10] C.S. Peskin. The immersed boundary method. *Acta Numerica*, 11:479–517, 2002.
- [11] F. Balboa Usabiaga, I. Pagonabarraga, and R. Delgado-Buscalioni. Inertial coupling for point particle fluctuating hydrodynamics. *J. Comp. Phys.*, 235:701–722, 2013.
- [12] F. Balboa Usabiaga, R. Delgado-Buscalioni, B. E. Griffith, and A. Donev. Inertial Coupling Method for particles in an incompressible fluctuating fluid. Submitted, code available at <https://code.google.com/p/fluum>, 2013.
- [13] A. Vázquez-Quesada, M. Ellero, and P. Español. Consistent scaling of thermal fluctuations in smoothed dissipative particle dynamics. *J. Chem. Phys.*, 130:034901, 2009.
- [14] A. Sierou and J. F. Brady. Accelerated Stokesian Dynamics simulations. *J. Fluid Mech.*, 448:115–146, 2001.
- [15] I. Theurkauff, C. Cottin-Bizonne, J. Palacci, C. Ybert, and L. Bocquet. Dynamic clustering in active colloidal suspensions with chemical signaling. *Phys. Rev. Lett.*, 108:268303, 2012.
- [16] B.E. Griffith, R.D. Hornung, D.M. McQueen, and C.S. Peskin. An adaptive, formally second order accurate version of the immersed boundary method. *J. Comput. Phys.*, 223(1):10–49, 2007. Software available at <http://ibamr.googlecode.com>.
- [17] Shih-Yuan Lu. Diffusion and reaction in regular arrays of spheres. *J. Chem. Phys.*, 109:4985, 1998.

- [18] Shih-Yuan Lu and Yi-Ming Yen. A first-passage scheme for determination of overall rate constants for non-diffusion-limited suspensions. *J. Chem. Phys.*, 116:3128, 2002.
- [19] M. R. Maxey and B. K. Patel. Localized force representations for particles sedimenting in Stokes flow. *International journal of multiphase flow*, 27(9):1603–1626, 2001.
- [20] Alexandre M Roma, Charles S Peskin, and Marsha J Berger. An adaptive version of the immersed boundary method. *J. Comput. Phys.*, 153(2):509–534, 1999.
- [21] Ricardo Cortez, Lisa Fauci, and Alexei Medovikov. The method of regularized Stokeslets in three dimensions: analysis, validation, and application to helical swimming. *Physics of Fluids*, 17:031504, 2005.
- [22] J. P. Hernandez-Ortiz, J. J. de Pablo, and M. D. Graham. Fast Computation of Many-Particle Hydrodynamic and Electrostatic Interactions in a Confined Geometry. *Phys. Rev. Lett.*, 98(14):140602, 2007.
- [23] Robert D Guy and Bobby Philip. A multigrid method for a model of the implicit immersed boundary equations. *Communications in Computational Physics*, 12(2):378, 2012.
- [24] B. Phillip R.D. Guy and B.E. Griffith. Geometric multigrid for an implicit-time immersed boundary method. Preprint, 2013.
- [25] Alan McKenney, Leslie Greengard, and Anita Mayo. A fast poisson solver for complex geometries. *Journal of Computational Physics*, 118(2):348–355, 1995.
- [26] T.T. Bringley and C.S. Peskin. Validation of a simple method for representing spheres and slender bodies in an immersed boundary method for Stokes flow on an unbounded domain. *J. Comp. Phys.*, 227(11):5397–5425, 2008.
- [27] David C Torney and Byron Goldstein. Rates of diffusion-limited reaction in periodic systems. *J. Stat. Phys.*, 49(3-4):725–750, 1987.
- [28] M. Skoge, A. Donev, F. H. Stillinger, and S. Torquato. Packing Hyperspheres in High-Dimensional Euclidean Spaces. *Phys. Rev. E*, 74:041127, 2006.
- [29] A. Donev, S. Torquato, and F. H. Stillinger. Pair Correlation Function Characteristics of Nearly Jammed Disordered and Ordered Hard-Sphere Packings. *Phys. Rev. E*, 71:011105, 2005.
- [30] Shih-Yuan Lu. Rate constants of spherical dispersions: From diffusion-limited data to nondiffusion limited results. *J. Chem. Phys.*, 110:12263, 1999.
- [31] James W Swan, John F Brady, Rachel S Moore, et al. Modeling hydrodynamic self-propulsion with stokesian dynamics. or teaching stokesian dynamics to swim. *Physics of Fluids*, 23:071901, 2011.
- [32] S. Lomholt and M.R. Maxey. Force-coupling method for particulate two-phase flow: Stokes flow. *J. Comp. Phys.*, 184(2):381–405, 2003.