

UC Davis

UC Davis Previously Published Works

Title

Scalable smoothing strategies for a geometric multigrid method for the immersed boundary equations

Permalink

<https://escholarship.org/uc/item/7qf470j6>

Authors

Bhalla, Amneet Pal Singh

Knepley, Matthew G

Adams, Mark F

et al.

Publication Date

2016-12-07

Peer reviewed

Scalable smoothing strategies for a geometric multigrid method for the immersed boundary equations

Amneet Pal Singh Bhalla · Matthew
G. Knepley · Mark F. Adams · Robert
D. Guy · Boyce E. Griffith

Received: date / Accepted: date

Abstract The immersed boundary (IB) method is a widely used approach to simulating fluid-structure interaction (FSI). Although explicit versions of

A.P.S.B. and B.E.G. acknowledge research funding from the National Institutes of Health (NIH award HL117063), the National Science Foundation (NSF awards ACI 1450327, CBET 1511427, and DMS 1410873), and The University of North Carolina at Chapel Hill. M.G.K. acknowledges research funding from the National Science Foundation (NSF award ACI 1450339) and the Department of Energy Office of Advanced Scientific Computing Research (U.S. DOE contract DE-AC02-06CH11357).

Amneet Pal Singh Bhalla
Department of Mathematics, University of North Carolina, Chapel Hill, NC
Carolina Center for Interdisciplinary Applied Mathematics, University of North Carolina,
Chapel Hill, NC
Tel.: +1-919-962-1294
Fax: +1-919-962-2568
E-mail: amneet@unc.edu

Matthew G. Knepley
Department of Computational and Applied Mathematics, Rice University, Houston, TX

Mark F. Adams
Scalable Solvers Group, Lawrence Berkeley National Laboratory, Berkeley, CA

Robert D. Guy
Department of Mathematics, University of California, Davis, CA

Boyce E. Griffith
Department of Mathematics, University of North Carolina, Chapel Hill, NC
Carolina Center for Interdisciplinary Applied Mathematics, University of North Carolina,
Chapel Hill, NC
Department of Biomedical Engineering, University of North Carolina, Chapel Hill, NC
McAllister Heart Institute, University of North Carolina, Chapel Hill, NC
Tel.: +1-919-962-1294
Fax: +1-919-962-2568
E-mail: boyceg@unc.edu

the IB method can suffer from severe time step size restrictions, these methods remain popular because of their simplicity and generality. In prior work [27], some of us developed a geometric multigrid preconditioner for a stable semi-implicit IB method under Stokes flow conditions; however, this solver methodology used a Vanka-type smoother that presented limited opportunities for parallelization. This work extends this Stokes-IB solver methodology by developing smoothing techniques that are suitable for parallel implementation. Specifically, we demonstrate that an additive version of the Vanka smoother can yield an effective multigrid preconditioner for the Stokes-IB equations, and we introduce an efficient Schur complement-based smoother that is also shown to be effective for the Stokes-IB equations. We investigate the performance of these solvers for a broad range of material stiffnesses, both for Stokes flows and flows at nonzero Reynolds numbers, and for thick and thin structural models. We show here that linear solver performance degrades with increasing Reynolds number and material stiffness, especially for thin interface cases. Nonetheless, the proposed approaches promise to yield effective solution algorithms, especially at lower Reynolds numbers and at modest-to-high elastic stiffnesses.

Keywords computational fluid dynamics · fluid-structure interaction · immersed boundary method · implicit time stepping · multigrid · scalability

Mathematics Subject Classification (2000) 65F08 · 65M55 · 76M20

1 Introduction

Since its introduction by Peskin [41, 42] to model blood flow through heart valves, the immersed boundary (IB) method has become a widely used approach to simulate fluid-structure interaction (FSI) in a broad range of scientific and engineering applications [43]. The flexibility of the IB approach to FSI has led to the development of many extensions, such as the ghost-cell IB method [50], the fictitious domain method [20], the immersed finite element method [54, 9], direct-forcing IB methods [51, 6], and an IB method for immersed reactive particles [8], along with other methods designed for various applications [32, 38, 10, 7, 48, 18]. Many of these methods use approaches that are rooted in Peskin’s original IB method, and modern versions of this method continue to see wide use, especially in biological applications.

A key feature of the IB approach to FSI is that it avoids mesh-conforming discretizations. Instead, the IB formulation of FSI uses a single momentum equation for both the fluid and the solid, which is expressed in Eulerian form, along with a Lagrangian description of the structural deformations and resulting forces. The Eulerian equations are discretized on a Cartesian grid, and the Lagrangian equations are approximated on a curvilinear mesh. Interaction between Eulerian and Lagrangian variables is mediated by discretized integral transforms with regularized delta function kernels. These transforms *interpolate* the Eulerian velocity onto the curvilinear mesh and *spread* the structural force density to the Eulerian grid.

There has been substantial work on both explicit and implicit versions of the IB method. A simple version of an explicit IB time stepping scheme first uses the current configuration of the structure to evaluate the structural forces; then spreads those forces to the Cartesian grid; solves the incompressible Navier-Stokes equations; interpolates the velocities back to the structure; and finally updates the configuration of the Lagrangian mesh using the interpolated velocity field. It is straightforward to develop more sophisticated versions of this method, e.g. that use Runge-Kutta schemes to increase the order of accuracy of the time discretization [34, 25]. The appeal of this explicit approach is that it requires a solver only for a Cartesian grid discretization of the incompressible Navier-Stokes equations. Fast solvers are readily available for these equations, including approaches based on fast Fourier transforms (FFT) for periodic domains and uniform Cartesian grids, or geometric multi-grid (MG) algorithms for other types of boundary conditions or locally refined Cartesian grids. Explicit time stepping is extremely effective for soft materials, but as the material stiffness α increases, the explicit treatment of the elastic forces imposes a time step size restriction, so that $\Delta t \sim \alpha^{-1}$ (for Stokes) or $\Delta t \sim \alpha^{-\frac{1}{2}}$ (for Navier-Stokes). Although it is not straightforward to analyze the Navier-Stokes case, a simple scaling argument implies that explicit IB method in Stokes flow conditions requires $\Delta t \sim \Delta x$ for thin elastic membranes and $\Delta t \sim \Delta x^3$ for thin beams. The stability restrictions for thick structures are less severe.

The alternatives to explicit IB methods are fully implicit and semi-implicit IB methods. Fully implicit IB methods can allow for the stable use of any time step size [40]. It is also possible to develop stable semi-implicit IB methods that use spreading and interpolation operators defined with respect to the current structural configuration, or an estimate of the new position [40]. These are referred to as *lagged* IB coupling operators [37, 14]. In effect, this approach linearizes the geometrical nonlinearities associated with the coupling operators. The structural configuration used to evaluate the Lagrangian forces still must be treated implicitly in such discretizations to maintain energy stability [40], but the resulting system of equations is substantially simpler than that of a fully implicit formulation.

Work on implicit IB formulations dates back to the first IB methods [41, 42], but here we briefly review research over only the past decade to develop efficient (semi-)implicit IB methods. Hou and Shi [29, 30] proposed a semi-implicit and unconditionally stable discretization of the IB equations for steady and unsteady Stokes flow for simple periodic interfaces with linear elasticity. They deploy a small-scale decomposition to obtain a formulation that can be expressed explicitly using Fourier transforms, which allows them to obtain an efficient solution method. Linear solvers based on semi-implicit discretizations of the IB method have also been proposed to treat more general structural geometries. Two notable examples are the works of Mori and Peskin [39] and Ceniceros et al. [14]. Both of these studies reformulate the IB equations by eliminating the Eulerian variables, so that the systems to be solved involve only Lagrangian degrees of freedom. An update to the Eule-

rian velocity and pressure is made thereafter by using the new position of the immersed structure. Mori and Peskin [39] suggest a simple diagonal preconditioner for the unstructured Lagrangian system, whereas Cenicerros et al. [14] employ an algebraic multigrid solver on the unstructured Lagrangian mesh, in which coarser and finer Lagrangian meshes are obtained by adding and removing Lagrangian points from a base mesh. Cenicerros et al. also advocate precomputing an explicit matrix-based representation of the Lagrangian linear operator for modest ratios of Lagrangian to Eulerian degrees of freedom. For periodic domains, they are able to apply the Lagrangian operator efficiently by exploiting the approximate translational invariance of Peskin’s regularized delta functions [43]. The idea of precomputing the Lagrangian matrix operator has also recently been employed by Kallemov et al. [31] and Usabiaga et al. [52] for rigid-body IB methods.

Although solving implicit or semi-implicit IB formulations using only Lagrangian variables can be very efficient for certain problems, developing scalable general-purpose algorithms for these formulations is difficult. In particular, constructing multigrid methods for Lagrangian formulations of the IB equations is challenging because the systems to be solved fundamentally involve the solution operator for the Stokes equations, which is used in this formulation to eliminate the Eulerian velocity and pressure variables. To avoid this difficulty, Guy et al. [27, 26] and Zhang et al. [55] proposed multigrid preconditioners for semi-implicit IB formulations in which the Lagrangian variables are eliminated. This approach requires the solution of Stokes-like systems of equations on structured Cartesian grids that involve only the Eulerian variables, thereby facilitating the development of geometric multigrid algorithms. Specifically, Guy et al. [27, 26] developed a geometric multigrid method for this Eulerian IB formulation similar to Vanka’s method for the Stokes equations [53], whereas Zhang et al. [55] proposed an approximate block-factorization preconditioner for this system that is similar to block multigrid preconditioners for the Navier-Stokes equations [17, 21, 12]. The approach of Zhang et al. does not appear to provide a robust semi-implicit solution strategy, and multiplicative smoothing strategies like that developed by Guy et al. present limited opportunities for parallelization. Further, Vanka-like smoothing for the IB equations requires the solution of relatively large block systems, resulting in a computational cost much greater than multigrid algorithms that can rely on simpler point relaxation smoothers, such as Jacobi or Gauss-Seidel smoothing.

This work extends the multigrid approach of Guy et al. [27] by introducing two different smoothing approaches that are amenable to large-scale parallelization. One smoother that we consider is similar to the Vanka-like scheme developed by Guy et al., but it uses a restricted additive Schwarz (RAS) method [13, 16] to couple the “big box” solves [27] required by this smoothing algorithm instead of the multiplicative algorithm developed by Guy et al. This approach allows each of the subdomain solves to be processed independently. We also develop a Schur complement (SC) smoother for the Stokes-IB equations based on an approximate block factorization. This approach is similar

to the method of Zhang et al. [55], except that here we use the block factorization as a smoother. We show that this SC smoother can be effective even when using only lightweight subdomain solves involving a few iterations of Chebyshev-accelerated Gauss-Seidel applied to Poisson-like operators.

Unlike the work of Guy et al. [27], here we consider nonzero Reynolds number flows in addition to the Stokes flow regime. As in earlier studies, solver convergence rates are shown to degrade with increasing material stiffness. This study also reveals, for the first time, that the linear solver convergence rates degrade with increasing Reynolds numbers. For low Reynolds numbers or Stokes flows, however, mesh-refinement studies demonstrate essentially optimal scaling under only a mild CFL-type time step size restriction for a range of material stiffnesses.

2 Immersed boundary method

2.1 Continuous equations of motion

In the immersed boundary (IB) formulation of fluid-structure interaction (FSI) problems, an Eulerian description is used for the momentum equation and divergence-free condition of both the fluid and the structure, and a Lagrangian description is used for the structural deformations and the resulting structural forces. We denote by $\mathbf{x} = (x_1, \dots, x_d) \in \Omega$ fixed Cartesian coordinates, in which $\Omega \subset \mathbb{R}^d$ is the fixed domain occupied by the entire fluid-structure system in d spatial dimensions. We denote by $\mathbf{s} = (s_1, \dots, s_d) \in U$ the fixed material coordinate system attached to the structure, in which $U \subset \mathbb{R}^d$ is the Lagrangian curvilinear coordinate domain. The position of the immersed structure at time t is denoted $\mathbf{X}(\mathbf{s}, t) \in \Omega$. To simplify the implementation, we consider only thin (codimension-1) massless structures and thick (codimension-0) neutrally buoyant bodies. In the case of a thick immersed body, the fluid and structure share the same uniform mass density ρ , and we further assume that the structure is viscoelastic with the same dynamic viscosity μ as the fluid. The equations of motion of the coupled fluid-structure system are [43]

$$\rho \left(\frac{\partial \mathbf{u}}{\partial t}(\mathbf{x}, t) + \mathbf{u}(\mathbf{x}, t) \cdot \nabla \mathbf{u}(\mathbf{x}, t) \right) = -\nabla p(\mathbf{x}, t) + \mu \nabla^2 \mathbf{u}(\mathbf{x}, t) + \mathbf{f}(\mathbf{x}, t), \quad (1)$$

$$\nabla \cdot \mathbf{u}(\mathbf{x}, t) = \mathbf{0}, \quad (2)$$

$$\mathbf{f}(\mathbf{x}, t) = \int_{\Omega} \mathbf{F}(\mathbf{s}, t) \delta(\mathbf{x} - \mathbf{X}(\mathbf{s}, t)) \, d\mathbf{s}, \quad (3)$$

$$\frac{\partial \mathbf{X}}{\partial t}(\mathbf{s}, t) = \int_U \mathbf{u}(\mathbf{x}, t) \delta(\mathbf{x} - \mathbf{X}(\mathbf{s}, t)) \, d\mathbf{x}, \quad (4)$$

$$\mathbf{F}(\mathbf{s}, t) = \mathcal{F}[\mathbf{X}(\cdot, t)](\mathbf{s}, t). \quad (5)$$

Eqs. (1) and (2) are the incompressible Navier-Stokes equations written in Eulerian form, in which $\mathbf{u}(\mathbf{x}, t)$ is the velocity, $p(\mathbf{x}, t)$ is the pressure, and $\mathbf{f}(\mathbf{x}, t)$ is the elastic force density. Eq. (5) determines the Lagrangian structural force

density from the configuration of the immersed structure via a functional $\mathcal{F} : \mathbf{X} \mapsto \mathbf{F}$. Interactions between Lagrangian and Eulerian quantities in Eqs. (3) and (4) are mediated by integral equations with Dirac delta function kernels, in which the d -dimensional delta function is $\delta(\mathbf{x}) = \prod_{i=1}^d \delta(x_i)$. Eq. (3) converts the Lagrangian force density $\mathbf{F}(\mathbf{s}, t)$ into an equivalent Eulerian density $\mathbf{f}(\mathbf{x}, t)$. The discretized version of this operation is called *force spreading*. We express the force spreading operation by $\mathbf{f} = \mathcal{S}[\mathbf{X}]\mathbf{F}$, in which $\mathcal{S}[\mathbf{X}]$ is the *force-spreading operator*. Eq. (4) determines the physical velocity of each Lagrangian material point from the Eulerian velocity field, so that the immersed structure moves according to the local value of the velocity field $\mathbf{u}(\mathbf{x}, t)$. This *velocity interpolation* operation is expressed as $\frac{\partial \mathbf{X}}{\partial t} = \mathcal{J}[\mathbf{X}]\mathbf{u}$, in which $\mathcal{J}[\mathbf{X}]$ is the *velocity-interpolation operator*. Notice that \mathcal{S} and \mathcal{J} are adjoint operators, $\mathcal{S} = \mathcal{J}^*$ [43].

2.2 Discrete equations of motion

We consider only linear solver performance in this work using linear systems of equations that arise from an energy-stable semi-implicit discretization of the IB equations. We use a spatial discretization that is similar to one used in earlier work [27], which is briefly described in A. To discretize these equations in time, let Δt be the time step size, and let n be the time step number. In each time step, we simultaneously solve for the updated Eulerian velocity \mathbf{u}^{n+1} and pressure p^{n+1} at time $t^{n+1} = (n+1)\Delta t$ along with the structural configuration \mathbf{X}^{n+1} . To simplify notation, we use $\mathcal{S}_h^n \equiv \mathcal{S}_h[\mathbf{X}^n]$ to indicate the spreading operator corresponding to structural configuration \mathbf{X}^n along with analogous notation for the interpolation operator \mathcal{J}_h . The time-stepping scheme reads

$$\rho \left(\frac{\mathbf{u}^{n+1} - \mathbf{u}^n}{\Delta t} + [\mathbf{u} \cdot \nabla_h \mathbf{u}]^{(n+\frac{1}{2})} \right) = -\nabla_h p^{n+1} + \mu \nabla_h^2 \mathbf{u}^{n+1} + \mathcal{S}_h^n \mathbf{F}^{n+1}, \quad (6)$$

$$\nabla_h \cdot \mathbf{u}^{n+1} = 0, \quad (7)$$

$$\frac{\mathbf{X}^{n+1} - \mathbf{X}^n}{\Delta t} = \mathcal{J}_h^n \mathbf{u}^{n+1}, \quad (8)$$

$$\mathbf{F}^{n+1} = \mathcal{F}_h[\mathbf{X}^{n+1}]. \quad (9)$$

Except for the nonlinear convection term, this scheme uses a combination of forward and backward Euler time stepping. In all tests reported herein, we omit the convection term, because including it does not affect linear solver performance. In applications, however, we often use a version of the PPM method [15, 44, 21] along with Adams-Bashforth to approximate the midstep value of $\mathbf{u} \cdot \nabla \mathbf{u}$ via

$$[\mathbf{u} \cdot \nabla_h \mathbf{u}]^{(n+\frac{1}{2})} = \frac{3}{2} \mathbf{u}^n \cdot \nabla_h \mathbf{u}^n - \frac{1}{2} \mathbf{u}^{n-1} \cdot \nabla_h \mathbf{u}^{n-1}. \quad (10)$$

Under reasonable assumptions on the form of the discretized force operator \mathcal{F}_h , the only stability restriction associated with this semi-implicit time stepping scheme is related to our explicit treatment of the convective term in the

momentum equation [40]. A fully implicit version of this scheme would replace $\mathcal{S}_h^n = \mathcal{S}_h[\mathbf{X}^n]$ by $\mathcal{S}_h^{n+1} = \mathcal{S}_h[\mathbf{X}^{n+1}]$, and likewise for \mathcal{J}_h . Such schemes do not appear to offer benefits in terms of energy stability or order of accuracy [40], but they do require the solution of a more complex system of nonlinear equations. The semi-implicit formulation used here can be seen as a method that linearizes the geometrical nonlinearities associated with the coupling operators \mathcal{S}_h and \mathcal{J}_h without sacrificing energy stability or formal order of accuracy.

Depending upon the functional form of the discrete force operator \mathcal{F}_h , Eqs. (6)–(9) can be linear or nonlinear. Because we consider only linear solver performance in this work, we use only linear force functionals of the form $\mathcal{F}_h[\mathbf{X}] = \mathcal{K}_h \mathbf{X}$, in which \mathcal{K}_h is the stiffness matrix of the elasticity model, and we solve Eqs. (6)–(9) by a preconditioned Krylov method. This requires the solution of linear systems of the form

$$\begin{pmatrix} \mathcal{A} & \mathcal{G} & -\mathcal{S}_h^n \mathcal{K}_h \\ -\mathcal{D} & \mathbf{0} & \mathbf{0} \\ -\mathcal{J}_h^n & \mathbf{0} & \frac{1}{\Delta t} \mathbf{I} \end{pmatrix} \begin{pmatrix} \mathbf{u}^{n+1} \\ p^{n+1} \\ \mathbf{X}^{n+1} \end{pmatrix} = \begin{pmatrix} \mathbf{g} \\ 0 \\ \frac{1}{\Delta t} \mathbf{X}^n \end{pmatrix}, \quad (11)$$

in which $\mathcal{A} = \frac{\rho}{\Delta t} \mathbf{I} - \mu \nabla_h^2$, $\mathcal{G} = \nabla_h$, and $\mathcal{D} = \nabla_h \cdot$ are block Eulerian operators, and \mathbf{g} contains contributions from previous time steps and explicitly evaluated terms from the current time step. To develop a system of equations amenable to solution via geometric multigrid methods, we use Eq. (8) to eliminate \mathbf{X}^{n+1} from the block system and obtain

$$\underbrace{\begin{pmatrix} \mathcal{A}_{\text{IB}} & \mathcal{G} \\ -\mathcal{D} & \mathbf{0} \end{pmatrix}}_{\mathcal{L}_{\text{IB}}} \begin{pmatrix} \mathbf{u}^{n+1} \\ p^{n+1} \end{pmatrix} = \begin{pmatrix} \mathbf{g} + \mathcal{S}_h^n \mathcal{K}_h \mathbf{X}^n \\ \mathbf{0} \end{pmatrix}, \quad (12)$$

in which $\mathcal{A}_{\text{IB}} = \mathcal{A} - \Delta t \mathcal{S}_h^n \mathcal{K}_h \mathcal{J}_h^n$ is the modified momentum operator that includes the projection of \mathcal{K}_h , the linear Lagrangian elasticity operator, onto the Eulerian frame. We refer to \mathcal{L}_{IB} as the *Stokes-IB operator*.

3 Multigrid

3.1 Basic multigrid algorithm

An effective preconditioner is needed to solve the Stokes-IB system (12) efficiently using a Krylov method. Here, we briefly discuss the key ingredients of a geometric multigrid (GMG) preconditioner for the Stokes-IB system. Detailed descriptions of the multigrid method are available [11, 49], and the development of GMG methods for the IB method is also discussed in previous work [27, 26].

We construct a hierarchy of uniform Cartesian discretizations of the spatial domain Ω . Let Ω^ℓ indicate a particular discretization with grid spacing h^ℓ , in which $\ell = 0, 1, \dots, \ell_{\text{max}}$ indicates the *level* of the discretization, with $\ell = 0$ denoting the coarsest level in the hierarchy and $\ell = \ell_{\text{max}}$ denoting the finest

Algorithm 1 recursive V-cycle multigrid

```

1: procedure  $\mathbf{w}^\ell \leftarrow \mathbf{MG}(\mathbf{w}^\ell, \mathbf{b}^\ell, \Omega^\ell, \nu_1, \nu_2)$ 
2:   if  $\Omega^\ell$  is the coarsest grid then
3:      $\mathbf{w}^\ell \leftarrow (\mathcal{L}_{\text{IB}}^\ell)^{-1} \mathbf{b}^\ell$  ▷ solve the coarse grid equation
4:   else
5:      $\mathbf{w}^\ell \leftarrow \mathbf{smooth}(\mathbf{w}^\ell, \mathbf{b}^\ell, \nu_1)$  ▷ apply  $\nu_1$  pre-smoothing sweeps
6:      $\mathbf{r}^\ell \leftarrow \mathbf{b}^\ell - \mathcal{L}_{\text{IB}}^\ell \mathbf{w}^\ell$  ▷ compute the residual on the present level
7:      $\mathbf{r}^{\ell-1} \leftarrow \mathcal{R}_\ell^{\ell-1} \mathbf{r}^\ell$  ▷ restrict the residual to the next coarser level
8:      $\mathbf{e}^{\ell-1} \leftarrow \mathbf{MG}(\mathbf{0}, \mathbf{r}^{\ell-1}, \Omega^{\ell-1}, \nu_1, \nu_2)$  ▷ recursively call MG
9:      $\mathbf{e}^\ell \leftarrow \mathcal{P}_{\ell-1}^\ell \mathbf{e}^{\ell-1}$  ▷ prolong the error from the next coarser level
10:     $\mathbf{w}^\ell \leftarrow \mathbf{w}^\ell + \mathbf{e}^\ell$  ▷ correct the solution on the present level
11:     $\mathbf{w}^\ell \leftarrow \mathbf{smooth}(\mathbf{w}^\ell, \mathbf{b}^\ell, \nu_2)$  ▷ apply  $\nu_2$  post-smoothing sweeps

```

level. The grid spacings on adjacent levels ℓ and $\ell - 1$ are related by an integer refinement ratio r_{ref} , so that $h^\ell = \frac{h^{\ell-1}}{r_{\text{ref}}}$. Here, we only consider $r_{\text{ref}} = 2$.

The basic V-cycle multigrid algorithm used in this work is shown in Algorithm 1. It aims to solve the discretized equations on Ω^ℓ by combining simple approximate solvers on level ℓ with coarse-grid corrections recursively computed on levels $\ell - 1, \ell - 2, \dots$. Specifically, on each grid level $\ell > 0$, Algorithm 1 uses a *smoother* to eliminate the high-frequency components of the error. The remaining low-frequency errors are meant to be eliminated by coarse-grid corrections. (Although not shown here, alternative multigrid algorithms, such as F- or W-cycles, seem to offer little to no benefit for the Stokes-IB equations with our present smoothers.)

Grid levels are connected by a *restriction operator* $\mathcal{R}_\ell^{\ell-1}$ that coarsens solution data from a finer level ℓ to a coarser level $\ell - 1$, and a *prolongation operator* $\mathcal{P}_{\ell-1}^\ell$ that interpolates solution data from a coarser level $\ell - 1$ to a finer level ℓ . These operators are also used to define the Stokes-IB system on coarser grid levels. Omitting the dependence on the time step number n , the block linear system on level ℓ is

$$\begin{pmatrix} \mathcal{A}^\ell - \Delta t [\mathcal{S}_h \mathcal{K}_h \mathcal{J}_h]^\ell & \mathcal{G}^\ell \\ -\mathcal{D}^\ell & \mathbf{0} \end{pmatrix} \begin{pmatrix} \mathbf{u}^\ell \\ p^\ell \end{pmatrix} = \begin{pmatrix} \mathbf{g}^\ell \\ \mathbf{0} \end{pmatrix}. \quad (13)$$

The coarse-grid Eulerian elasticity operator is defined for $\ell < \ell_{\text{max}}$ via a Galerkin projection,

$$[\mathcal{S}_h \mathcal{K}_h \mathcal{J}_h]^\ell = (\mathcal{R}_\mathbf{u})_{\ell+1}^\ell [\mathcal{S}_h \mathcal{K}_h \mathcal{J}_h]^{\ell+1} (\mathcal{P}_\mathbf{u})_\ell^{\ell+1}, \quad (14)$$

in which $\mathcal{R}_\mathbf{u}$ and $\mathcal{P}_\mathbf{u}$ are restriction and prolongation operators for velocity-like degrees of freedom only. The coarse-grid operators \mathcal{A}^ℓ , \mathcal{G}^ℓ , and \mathcal{D}^ℓ are constructed by rediscrretization. We denote the block system (13) by

$$\mathcal{L}_{\text{IB}}^\ell \mathbf{w}^\ell = \mathbf{b}^\ell. \quad (15)$$

Given an approximate solution $\tilde{\mathbf{w}}^\ell$, the corresponding error equation is

$$\mathcal{L}_{\text{IB}}^\ell \mathbf{e}^\ell = \mathbf{r}^\ell, \quad (16)$$

in which $\mathbf{e}^\ell = \mathbf{w}^\ell - \tilde{\mathbf{w}}^\ell$ is the error and $\mathbf{r}^\ell = \mathbf{b}^\ell - \mathcal{L}_{\text{IB}}^\ell \tilde{\mathbf{w}}^\ell$ is the residual.

3.2 Prolongation and restriction

The velocity prolongation operator $\mathcal{P}_{\mathbf{u}}$ is based on lowest-order Raviart-Thomas interpolation on quadrilaterals [2], which uses piecewise-linear interpolation in the normal direction to the cell face along with piecewise-constant interpolation in the tangential direction. The velocity restriction operator is taken to be the adjoint of the prolongation operator, $\mathcal{R}_{\mathbf{u}} = \mathcal{P}_{\mathbf{u}}^*$. Consequently, the coarse-grid versions of the Eulerian elasticity operator $\mathcal{S}_h \mathcal{K}_h \mathcal{J}_h$ retain the symmetry of the fine-grid operator. We also use linear interpolation to prolong pressure values from coarse to fine grid levels, and we use simple averaging to restrict pressures from fine to coarse levels. In this case $\mathcal{R}_p \neq \mathcal{P}_p^*$, but preliminary tests suggest that this has essentially no effect on solver convergence.

3.3 Smoothers

We consider two classes of smoothers: Vanka-like algorithms based on Schwarz domain decomposition methods, and algorithms based on an approximate block factorization. In both cases, the basic algorithm may not act as a smoother (i.e. it may fail to damp some high-frequency error components). To provide enhanced smoothing without introducing a damping parameter, we use a fixed number of flexible GMRES (FGMRES) [45] iterations preconditioned by the basic algorithm.

3.3.1 Schwarz smoothing

Additive and multiplicative Schwarz are domain decomposition methods [47] that solve restricted versions of the linear system on overlapping subdomains. The key difference between additive and multiplicative domain decomposition methods is that, in an additive method, the subdomain-restricted equations are solved independently, whereas in a multiplicative method, the most recently computed solution values from all subdomains are used in computing the residual for each subdomain solve. Additive and multiplicative Schwarz are thereby generalizations of the classical Jacobi and Gauss-Seidel methods. Multiplicative algorithms are fundamentally sequential, making them difficult or impossible to parallelize. Additive algorithms, by contrast, are readily parallelized.

In this work, we consider two Schwarz-like algorithms. One is *restricted additive Schwarz (RAS)* [13, 16] (Algorithm 2) with overlapping subdomains that correspond to all of the degrees of freedom associated with contiguous, rectangular boxes of grid cells. The other is a multiplicative version of the RAS algorithm, which we refer to as *restricted multiplicative Schwarz (RMS)* (Algorithm 3). Notice that the only difference between the two algorithms is the manner in which the residual is computed for each of the subdomain solves. Our RMS smoother is similar to the “big-box” Vanka algorithm described by

Algorithm 2 restricted additive Schwarz (RAS)

```

1: procedure  $\mathbf{w} \leftarrow \mathbf{RAS}(\mathbf{w}, \mathbf{b})$ 
2:   partition  $\Omega$  into  $N$  overlapping subdomains  $\Omega_i$  with index sets  $G_i$  and restriction
   matrices  $R_i$ , and construct non-overlapping index sets  $\tilde{G}_i \subset G_i$  and corresponding
   restriction matrices  $\tilde{R}_i$ 
3:    $\mathbf{r} \leftarrow \mathbf{b} - \mathcal{L}\mathbf{w}$  ▷ form the residual for the initial value of  $\mathbf{w}$ 
4:   for partition  $i = 1 \dots N$  do
5:      $\mathcal{L}_i \leftarrow R_i \mathcal{L} R_i^T$  ▷ construct the subdomain operator  $\mathcal{L}_i$ 
6:      $\mathbf{b}_i \leftarrow R_i \mathbf{r}$  ▷ form the local right-hand side  $\mathbf{b}_i$ 
7:      $\mathbf{w} \leftarrow \mathbf{w} + \tilde{R}_i^T (\mathcal{L}_i)^{-1} \mathbf{b}_i$  ▷ perform a local solve and update  $\mathbf{w}$  in  $\tilde{G}_i$ 

```

Algorithm 3 restricted multiplicative Schwarz (RMS)

```

1: procedure  $\mathbf{w} \leftarrow \mathbf{RMS}(\mathbf{w}, \mathbf{b})$ 
2:   partition  $\Omega$  into  $N$  overlapping subdomains  $\Omega_i$  with index sets  $G_i$  and restriction
   matrices  $R_i$ , and construct non-overlapping index sets  $\tilde{G}_i \subset G_i$  and corresponding
   restriction matrices  $\tilde{R}_i$ 
3:   for partition  $i = 1 \dots N$  do
4:      $\mathbf{r} \leftarrow \mathbf{b} - \mathcal{L}\mathbf{w}$  ▷ form the residual for the updated value of  $\mathbf{w}$ 
5:      $\mathcal{L}_i \leftarrow R_i \mathcal{L} R_i^T$  ▷ construct the subdomain operator  $\mathcal{L}_i$ 
6:      $\mathbf{b}_i \leftarrow R_i \mathbf{r}$  ▷ form the local right-hand side  $\mathbf{b}_i$ 
7:      $\mathbf{w} \leftarrow \mathbf{w} + \tilde{R}_i^T (\mathcal{L}_i)^{-1} \mathbf{b}_i$  ▷ perform a local solve and update  $\mathbf{w}$  in  $\tilde{G}_i$ 

```

Guy et al. [27]. We consider the effect of different subdomain sizes and overlap widths on the performance of both algorithms.

3.3.2 Schur complement (SC) smoothing

An alternative smoothing approach is to construct an approximate block factorization of $(\mathcal{L}_{\text{IB}}^\ell)^{-1}$ that separates the block system into velocity and pressure subdomain operators. Omitting the superscript “ ℓ ” for notational clarity, we first note that $\mathcal{L}_{\text{IB}}^{-1}$ can be written as

$$\mathcal{L}_{\text{IB}}^{-1} = \begin{pmatrix} \mathbf{I} & -\mathcal{A}_{\text{IB}}^{-1} \mathcal{G} \\ \mathbf{0} & \mathbf{I} \end{pmatrix} \begin{pmatrix} \mathcal{A}_{\text{IB}}^{-1} & \mathbf{0} \\ \mathbf{0} & \mathcal{M}^{-1} \end{pmatrix} \begin{pmatrix} \mathbf{I} & \mathbf{0} \\ \mathcal{D} \mathcal{A}_{\text{IB}}^{-1} & \mathbf{I} \end{pmatrix}, \quad (17)$$

in which $\mathcal{M} = \mathcal{D} \mathcal{A}_{\text{IB}}^{-1} \mathcal{G}$ is the *Schur complement* of the block system of equations in (12). The cost of constructing the full factorization of $\mathcal{L}_{\text{IB}}^{-1}$ is prohibitive, but in practice, $\mathcal{L}_{\text{IB}}^{-1}$ rarely needs to be formed explicitly. Here, we only need to be able to apply this operator to known right-hand side vectors. Moreover, the application of the exact inverse operator can be unnecessary in a smoother, since we only wish to eliminate high-frequency error modes and not low-frequency modes. Consequently, in a multigrid algorithm, operators that approximate the action of $\mathcal{A}_{\text{IB}}^{-1}$ and \mathcal{M}^{-1} can suffice.

There are many choices for approximating $\mathcal{A}_{\text{IB}}^{-1}$ and \mathcal{M}^{-1} , but we find that simple approximate solvers for \mathcal{A}_{IB} and \mathcal{M} lead to an effective smoothing algorithm. We approximate the action of $\mathcal{A}_{\text{IB}}^{-1}$ by $\widetilde{\mathcal{A}}_{\text{IB}}^{-1}$, which uses a fixed number of Chebyshev iterations preconditioned by Gauss-Seidel applied to

\mathcal{A}_{IB} , and we approximate the action of \mathcal{M}^{-1} by $\widetilde{\mathcal{M}}^{-1}$, which uses a fixed number of Chebyshev iterations for the operator $(\mathcal{D}\widetilde{\mathcal{A}}_{IB}^{-1}\mathcal{G})$ preconditioned by Gauss-Seidel applied to a sparse approximate Schur complement, $\widehat{\mathcal{M}} = \mathcal{D}(\text{diag}(\mathcal{A}_{IB}))^{-1}\mathcal{G}$. Notice that the sparse approximate Schur complement $\widehat{\mathcal{M}}$ takes the form of an inhomogeneous discrete Poisson operator. We typically use two Chebyshev iterations for both $\widetilde{\mathcal{A}}_{IB}^{-1}$ and $\widetilde{\mathcal{M}}^{-1}$. With these operators so defined, we specify the action of $\widetilde{\mathcal{L}}_{IB}^{-1} \approx \mathcal{L}_{IB}^{-1}$ by

$$\widetilde{\mathcal{L}}_{IB}^{-1} = \begin{pmatrix} \mathbf{I} & -\widetilde{\mathcal{A}}_{IB}^{-1}\mathcal{G} \\ \mathbf{0} & \mathbf{I} \end{pmatrix} \begin{pmatrix} \widetilde{\mathcal{A}}_{IB}^{-1} & \mathbf{0} \\ \mathbf{0} & \widetilde{\mathcal{M}}^{-1} \end{pmatrix} \begin{pmatrix} \mathbf{I} & \mathbf{0} \\ \mathcal{D}\widetilde{\mathcal{A}}_{IB}^{-1} & \mathbf{I} \end{pmatrix}. \quad (18)$$

Notice that this smoother involves only point-relaxation. By contrast, in two spatial dimensions, the Vanka-like smoothers relax over Cartesian boxes of size $n_x \times n_y$, which requires solving linear systems with $O(n_x n_y)$ variables. Consequently, the typical cost of an application of the SC smoother is substantially less than the cost of one application of the Vanka-like smoother.

4 Software implementation

The solvers for this semi-implicit IB method are implemented in the open-source IBAMR library [1]. We use PETSc [5, 3, 4] to provide Krylov solvers and Schur complement-based smoothers. A custom implementation of the Schwarz algorithms is provided by IBAMR. IBAMR also relies on SAMRAI [28, 46] for Cartesian grid management, and for inter-level data transfer operations.

5 Results

Our tests explore the linear solver performance for both thick and thin elastic structures at a range for material stiffnesses and flow conditions. We consider immersed structures that are modeled using systems of elastic fibers. For simplicity, we assume that the leading Lagrangian coordinate s_1 varies along the direction of each fiber, and that the remaining curvilinear coordinates serve to label a particular fiber. The unit fiber tangent vector is $\boldsymbol{\tau} = \frac{\partial \mathbf{X}}{\partial s_1} / \left\| \frac{\partial \mathbf{X}}{\partial s_1} \right\|$, and the tension T in each fiber is a function of the fiber strain $\left\| \frac{\partial \mathbf{X}}{\partial s_1} \right\|$. It can be shown that the Lagrangian fiber force functional takes the form [43]

$$\mathbf{F} = \mathcal{F}[\mathbf{X}] = \frac{\partial(T\boldsymbol{\tau})}{\partial s_1}. \quad (19)$$

To explore the performance of the linear solvers, it is convenient to use fibers with zero resting lengths that resist only extension, so that the fiber tension is




Fig. 1: Representative results from the thick elastic test problem of Sec. 5.1. Here, we plot the velocity and pressure field along with the structure configuration for $\mu = 0$ and a relative stiffness of $\gamma = 5$. The Eulerian grid spacing is $h = \frac{1}{128}$. For clarity, only half of the Lagrangian fibers are shown. Results for different values of μ and γ are similar.

$T = \alpha \left\| \frac{\partial \mathbf{X}}{\partial s_1} \right\|$, in which α is the fiber stiffness. The resulting Lagrangian force density is

$$\mathbf{F} = \alpha \frac{\partial^2 \mathbf{X}}{\partial s_1^2}, \quad (20)$$

which is a linear functional. We use a simple second-order finite difference approximation to this functional, as described in A.2.

In all of our tests, the physical domain is $\Omega = [0, 1]^2$, and we use regularized lid-driven cavity flow boundary conditions, for which the velocity is set to zero along $\partial\Omega$ except along the top wall, where we prescribe $u(x, 1) = (1 - \cos(2\pi x))/2$ and $v = 0$. In these tests, the physical boundary conditions set a characteristic flow speed that determines the dynamic timescale of interest, independent of the elastic timescales of the immersed structure. The domain size sets the characteristic lengthscale as $L = 1$, and lid-driven cavity flow conditions set the characteristic flow speed as $U = 1$, so that $Re = \frac{\rho U L}{\mu} = \frac{\rho}{\mu}$. We set $\rho = 1$ for nonzero Reynolds number cases, yielding $\mu = Re^{-1}$.

We only consider square computational domains, and in our multigrid algorithm, we always use an 8×8 coarse grid along with sufficiently many finer levels to reach the targeted Eulerian grid spacing. A direct solver is used on the coarsest grid level. As mentioned previously, we only consider $r_{\text{ref}} = 2$.

5.1 Thick elastic shell

We first consider a thick circular annulus described using Lagrangian curvilinear coordinates $(s_1, s_2) \in U = [0, 2\pi) \times [0, w]$, in which $w = 1/16$ is the

thickness of the annulus. The structure is initially placed at the center of the domain in the configuration

$$\mathbf{X}(s_1, s_2) = (x_c + (r + s_2) \cos(s_1), y_c + (r + s_2) \sin(s_1)). \quad (21)$$

We choose the center to be $\mathbf{x}_c = (x_c, y_c) = (0.5, 0.5)$ and the inner radius to be $r = 1/4$. This configuration has been used as a standard test case in the IB literature [27, 9, 25, 24, 23]. The Eulerian domain is discretized using a uniform $N \times N$ grid, so that the Cartesian grid spacing is $h = \Delta x_1 = \Delta x_2 = 1/N$. For the Lagrangian domain, we use $M_1 = \frac{19}{8}N$ points in the s_1 direction and $M_2 = \frac{3}{32}N + 1$ points in the s_2 direction, which yields a physical spacing between the Lagrangian nodes approximately equal to $\frac{2}{3}h$. In these tests, we always use $\Delta t = 0.32\Delta x$. This implies that the time step size satisfies a mild advective CFL-type condition under grid refinement.

We characterize the stiffness α in terms of a stiffness ratio γ via

$$\alpha = \gamma \frac{3.93}{0.005}. \quad (22)$$

In prior work [27], $\alpha = \frac{3.93}{0.005}$ was shown to be approximately the largest stiffness for this problem for which the explicit solver is stable for a time step size of $\Delta t = 0.005$ and a grid spacing of $\Delta x = \frac{1}{64}$ in Stokes flow conditions. This largest stiffness is relatively insensitive to grid refinement [27]. Thus, γ roughly characterizes the ratio of the stiffness α and the maximum stiffness that can be used by an explicit time stepping scheme, α_{explicit} . For non-zero Reynolds numbers, α_{explicit} also depends on the fluid viscosity μ in a manner that we do not explore in this work. Thus, γ only *approximately* characterizes the ratio of the elastic stiffness to the largest stiffness that can be used by an explicit time stepping scheme. We consider a range of values of γ , from relatively soft ($\gamma = 5$) to very stiff ($\gamma = 500$). Representative results are shown in Fig. 1.

5.1.1 Schwarz smoothers

In these tests, we execute a single time step of the semi-implicit IB time integrator with the convective term disabled, so that we can focus on linear solver performance. Each linear solve is allowed to run until the initial residual is reduced by 10^{-12} or it reaches 100 iterations. We first examine the effect of subdomain size and overlap width on solver performance. For these tests, we use $N = 128$ and $\gamma = 500$, which is considered “very stiff” in our previous work [27], and consider Stokes flow (Fig. 2) and time-dependent flows for $\mu = 1.0, 0.1, 0.01$ (Fig. 3). We use subdomains of size $4 \times 4, 8 \times 8,$ and 16×16 and consider 0, 2, or 4 cells of overlap. It is clear that solver performance degrades substantially as the fluid viscosity decreases. At the highest Reynolds numbers, only the largest subdomains and overlaps yield effective solvers. For easier cases, the multiplicative and additive solvers yield similar performance, but in many of the more challenging cases, the multiplicative smoother can




Fig. 2: Multiplicative (solid line) and additive (dashed line) smoother performance under Stokes flow conditions for the thick elastic shell (Sec. 5.1) with $N = 128$ and $\gamma = 500$, and for different subdomain sizes and overlap widths. For this problem, solver performance is relatively insensitive to subdomain size, overlap width, or smoother algorithm.

converge in approximately half the iterations as the additive smoother for smaller overlap widths. At lower Reynolds numbers and Stokes flow conditions, however, the additive and multiplicative algorithms yield similar performance.

We also examine the scalability of the multiplicative and additive smoothers under grid refinement. Fig. 4 shows the performance of the additive and multiplicative smoothers using 8×8 subdomains with an overlap of 2 for various relative stiffnesses and viscosities, and Fig. 5 shows results from similar tests using 16×16 subdomains with an overlap of 4. It is clear that at low Reynolds numbers or low relative stiffnesses, the multiplicative smoother yields an essentially scalable algorithm, as shown previously [27]. At high elastic stiffnesses and larger Reynolds numbers, both solvers begin to break down, but with sufficiently large subdomain sizes and overlap widths, the additive version of the algorithm yields performance that is similar to that obtained by the multiplicative algorithm.

5.1.2 Schur complement smoother

As in the tests for the Schwarz preconditioners, we execute a single time step of the semi-implicit IB time integrator with the convective term disabled. We perform scalability tests using the Schur complement-based smoother for Stokes flows and for time-dependent flows with various viscosities at various relative elastic stiffnesses. Results are summarized in Fig. 6. The Schur complement-based smoother is more robust under both increasing elastic stiffness and decreasing viscosity than the Schwarz-based algorithms except for the largest elastic stiffnesses. Notice that the only case where the solver fails to reach its tight convergence threshold of 10^{-12} is for $\mu = 0.001$ and $\gamma = 500$. On the other hand, the Schur complement approach generally requires some-

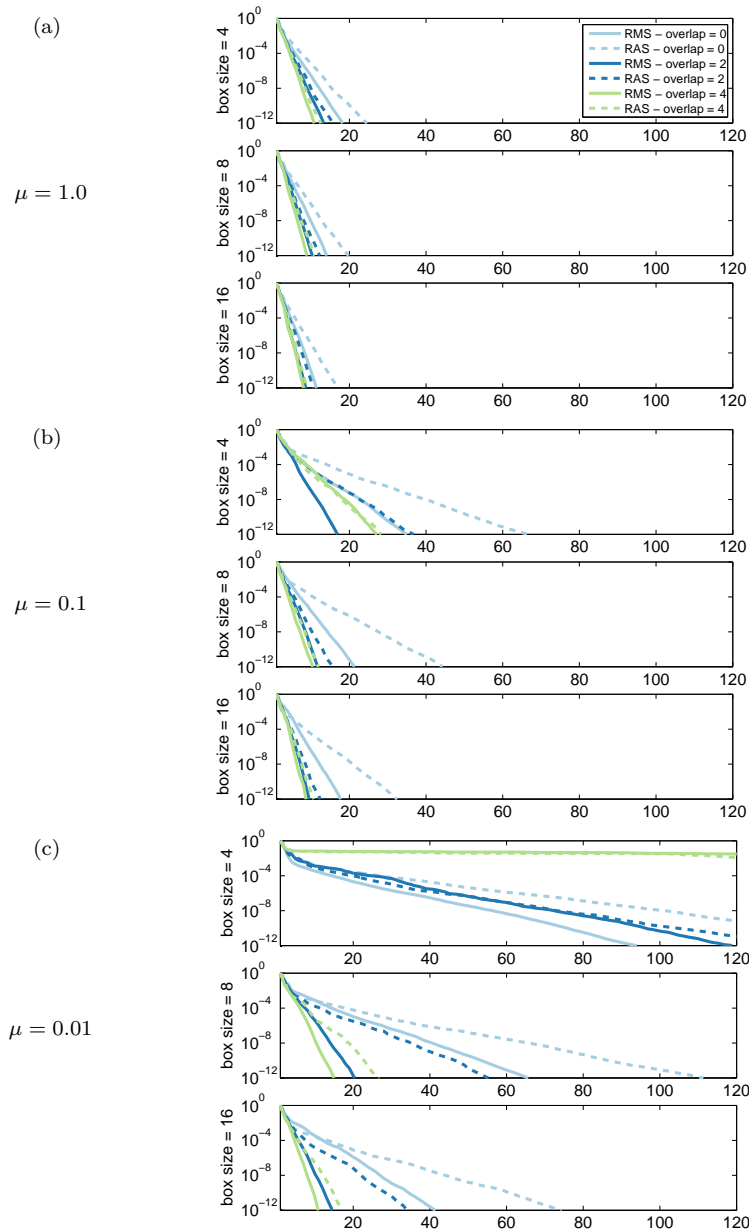


Fig. 3: Similar to Fig. 2, but here showing the effect of decreasing viscosity, with (a) $\mu = 1$, (b) $\mu = 0.1$, and (c) $\mu = 0.01$. Performance clearly degrades with decreasing viscosity. In the most challenging cases, the multiplicative smoother outperforms the additive smoother by a wide margin.




Fig. 4: Performance of multigrid using (a) multiplicative (RMS) and (b) additive (RAS) smoothers under grid refinement for a range of relative stiffnesses (γ) and viscosities (μ), using subdomains of size 8×8 and an overlap width of 2. In most cases, using the additive smoother results in only a modest increase in iterations compared to the multiplicative algorithm.

what more multigrid iterations than the Schwarz-based method for a given set of model parameters. However, each application of the SC smoother is substantially less expensive than the RAS/RMS smoothers, and in our current implementation, we typically find that the SC-based solver outperforms the RAS/RMS solver in terms of total wall-clock time.

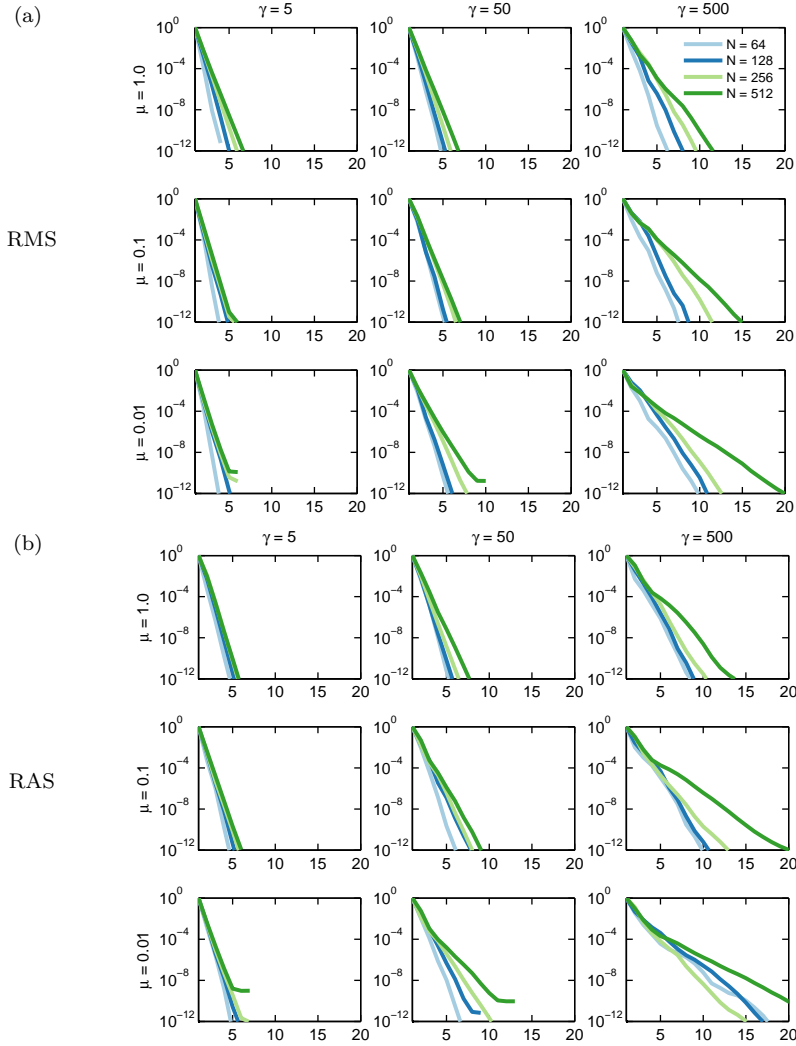


Fig. 5: Similar to Fig. 4, but here using subdomains of size 16×16 and an overlap width of 4. With larger subdomains and overlap widths, the performance of the additive algorithm approaches that of the multiplicative algorithm for a broader range of physical parameters (compare to Fig. 4).

5.2 Thin elastic membrane

Next, we consider a thin circular membrane described using Lagrangian curvilinear coordinates $s_1 \in U = [0, 2\pi)$ with initial configuration

$$\mathbf{X}(s_1) = (x_c + r \cos(s_1), y_c + r \sin(s_1)). \quad (23)$$

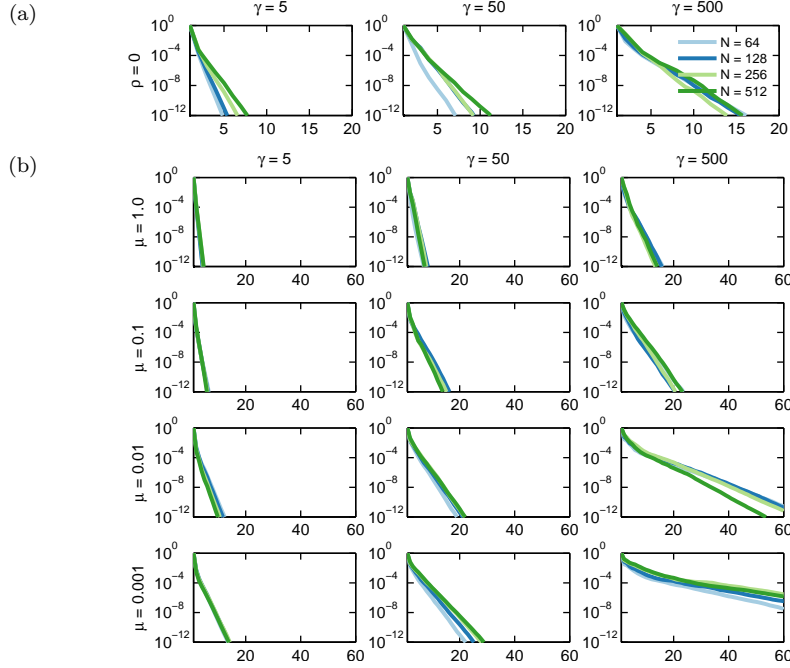


Fig. 6: Performance of the Schur complement-based smoother for the thick elastic shell (Sec. 5.1). We consider the effects of grid refinement for a range of relative stiffnesses (γ) and viscosities (μ), for both (a) Stokes flow conditions ($\rho = 0$) and (b) time-dependent flow conditions with decreasing amounts of fluid viscosity. Notice that the Schur complement-based smoother yields a more robust algorithm except for the highest stiffnesses at low viscosities; compare to Figs 4 and 5.

As in Sec. 5.1, we choose the center to be $\mathbf{x}_c = (x_c, y_c) = (0.5, 0.5)$ and the radius to be $r = 1/4$. We again use

$$\mathbf{F} = \alpha \frac{\partial^2 \mathbf{X}}{\partial s_1^2}, \quad (24)$$

and, as before, the Eulerian domain is discretized using an $N \times N$ grid, and the Lagrangian coordinates are discretized using $M_1 = \frac{19}{8}N$ points in the s_1 direction.

We again characterize the stiffness α in terms of a stiffness ratio γ , but α is now defined via

$$\alpha = 7\gamma \frac{3.93}{0.005}, \quad (25)$$

which yields approximately the same total force as in the thick interface case. In this thin case, at a fixed time step size, the maximum stiffness allowed by an explicit solver *decreases* in proportion to the grid spacing. This is in contrast to the thick case. Thus, for a fixed mechanical stiffness, the *numerical*




Fig. 7: Representative results from the thick elastic test problem of Sec. 5.2. Here, we plot the velocity and pressure field along with the structure configuration for $\mu = 0$ and a relative stiffness of $\gamma = 5$. The Eulerian grid spacing is $h = \frac{1}{128}$. Results for different values of μ and γ are similar.

stiffness of the problem increases under grid refinement. Consequently, the thin interface case poses substantially greater challenges to the solvers. Moreover, an analysis similar to that presented for a thick elastic shell [27] implies that $\gamma = 5$ is approximately a factor of 100 times stiffer than the largest elastic stiffness permitted by an explicit solver at $\Delta x = \frac{1}{64}$ and $\Delta t = 0.005$. Thus, the thin cases considered here are much more numerically challenging than the thick cases considered above. Representative results are shown in Fig. 7.

5.2.1 Schwarz smoothers

As before, we disable the convective term in our semi-implicit time integrator to focus on linear solver performance. Fig. 8 shows the effect of subdomain size and overlap width on solver performance for the multiplicative and additive smoothers in Stokes flow conditions at a relative stiffness of $\gamma = 500$. In this thin interface case, solver performance depends strongly on overlap width. This is in contrast to the case of a thick elastic shell (e.g. Fig. 2). At an overlap width of 4, the two Schwarz smoothers yield nearly identical convergence rates for all subdomain sizes considered in the Stokes flow case. Time dependent flow conditions pose a greater challenge to the solver (Fig. 9). At $\mu = 1$, solver performance is largely insensitive to subdomain size and overlap width, except for an overlap width of 0. By contrast, for $\mu = 0.1$, there are substantial differences in solver performance for the different subdomain sizes and overlap widths. For $\mu = 0.01$, only the largest subdomain sizes and overlap widths yield convergent solver algorithms. It is clear that the thin interface case is fundamentally *more stiff* than the thick body case.




Fig. 8: Multiplicative (solid line) and additive (dashed line) smoother performance under Stokes flow conditions for the thin elastic membrane (Sec. 5.2) with $N = 128$ and $\gamma = 500$, and for different subdomain sizes overlap widths. Unlike the case of a thick elastic shell, in this case solver performance has a strong dependence on overlap width. With an overlap of 4, the additive and multiplicative smoothers yield similar performance.

Fig. 10 shows the effects of grid refinement on solver performance for the time-dependent cases. For these tests, we consider only a subdomain size of 16×16 along with an overlap width of 4. The multiplicative smoother yields an essentially scalable multigrid algorithm except for the highest stiffness ($\gamma = 500$) and lowest viscosity ($\mu = 0.01$) considered. Performance of the additive smoother is similar except for $\gamma = 500$ and $\mu = 0.01$. In this challenging case, both smoothers show poor performance for $N = 256$, and the additive smoother stagnates at the highest grid spacing ($N = 512$). In practical time-dependent calculations, we likely would use a relative convergence tolerance around 10^{-6} , which corresponds to 4–5 multigrid iterations in all but the most difficult cases considered here.

5.2.2 Schur complement smoothers

As in the tests for the Schwarz preconditioners, we execute a single time step of the semi-implicit IB time integrator with the convective term disabled, now using the Schur complement-based smoother. Fig. 11 summarizes solver performance under grid refinement for a range of flow conditions and elastic stiffnesses. The Schur complement-based smoother is extremely robust for Stokes flows and low Reynolds number cases. However, the solver begins to stagnate for the higher stiffness cases even at a modest Reynolds number of 100. For this thin interface case, it appears that the current additive and multiplicative Schwarz smoothers are more effective, although they also struggle with higher Reynolds numbers and elastic stiffnesses.

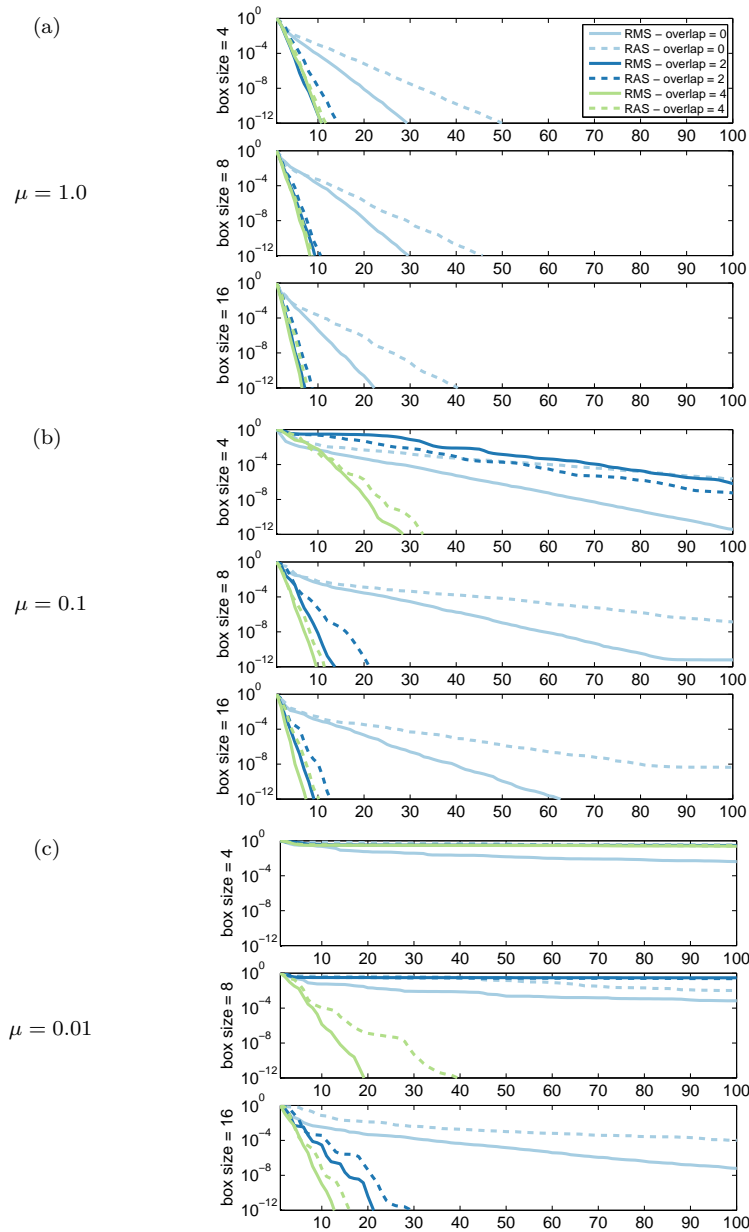


Fig. 9: Similar to Fig. 8, but here showing the effect of decreasing viscosity, with (a) $\mu = 1$, (b) $\mu = 0.1$, and (c) $\mu = 0.01$. As in the thick case (Sec. 5.1), performance degrades with decreasing viscosity, and the multiplicative smoother can outperform the additive smoother by a wide margin.

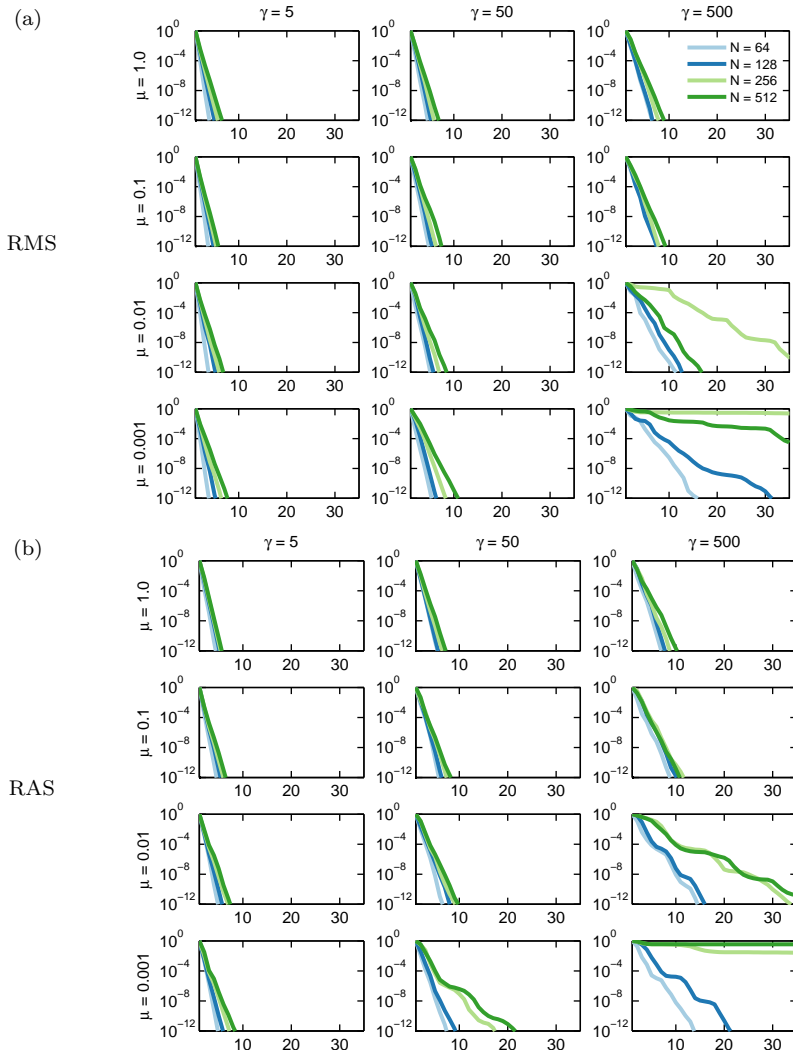


Fig. 10: Performance of multigrid using (a) multiplicative (RMS) and (b) additive (RAS) smoothers under grid refinement for a range of relative stiffnesses (γ) and viscosities (μ), using subdomains of size 16×16 and an overlap width of 4. Notice that the performance of the additive algorithm is similar to that of the multiplicative algorithm in most cases. Both algorithms ultimately stall for the largest stiffnesses for sufficiently small viscosity.

5.3 Suspension of immersed structures

This test case is similar to the thin interface case of Sec. 5.2, but here we consider a suspension of 16 structures, each with an initial configuration corresponding to a circle of radius $r = 1/16$. The structures are randomly placed

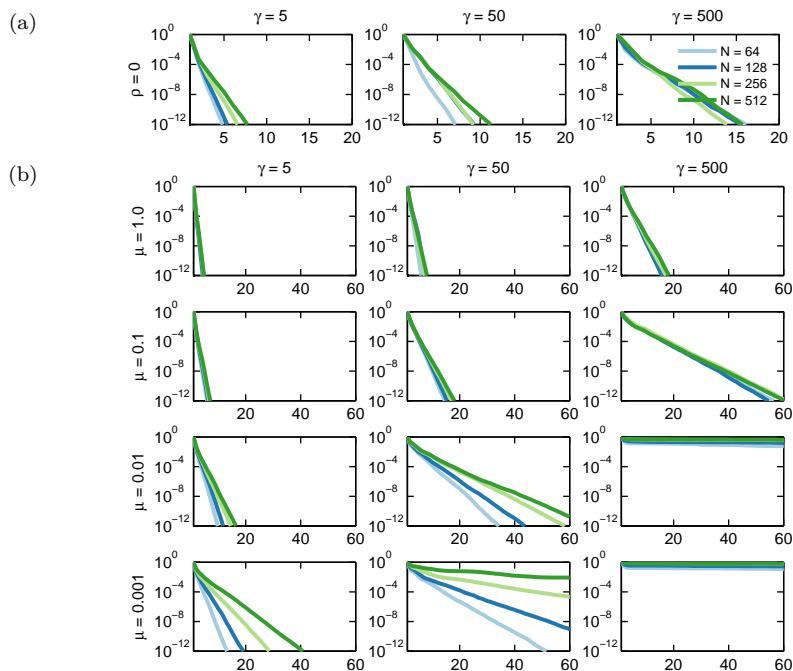


Fig. 11: Performance of the Schur complement-based smoother for the thin elastic membrane (Sec. 5.2). We consider the effects of grid refinement for a range of relative stiffnesses (γ) and viscosities (μ), for (a) Stokes flow conditions ($\rho = 0$) and (b) time-dependent flow conditions with decreasing amounts of fluid viscosity. The Schur complement-based smoother is very effective for Stokes flow conditions and lower viscosities, but at higher Reynolds numbers, performance degrades even at relatively small stiffnesses ($\gamma = 5$), and the solver stagnates for the higher stiffnesses considered.

in the domain and are required not to overlap each other or the domain boundary. Fig. 12 shows the distribution of structures along with the resulting pressure field for $\gamma = 5$. We consider only the RAS and Schur complement smoothers, and we explore the performance of the solver with increasing numbers of processors for $N = 64, 128, 256,$ and 512 , using $(N/64)^2$ processors for each case, so that the number of grid cells assigned to each processor remains fixed. Fig. 12 shows the Cartesian grid-based parallel domain decompositions. Fig. 13 summarizes the solver performance under grid refinement for a range of flow conditions and elastic stiffnesses. Performance is similar to that obtained in serial for the case of a single immersed membrane, although the Schur complement-based smoother shows slightly poorer scaling in parallel than in serial. This is not unexpected because the Schur complement-based solver uses processor-restricted Gauss-Seidel in its subdomain operators rather than a true parallel Gauss-Seidel algorithm. Nonetheless, the Schur complement-based solver yields good scalability in cases where the underlying

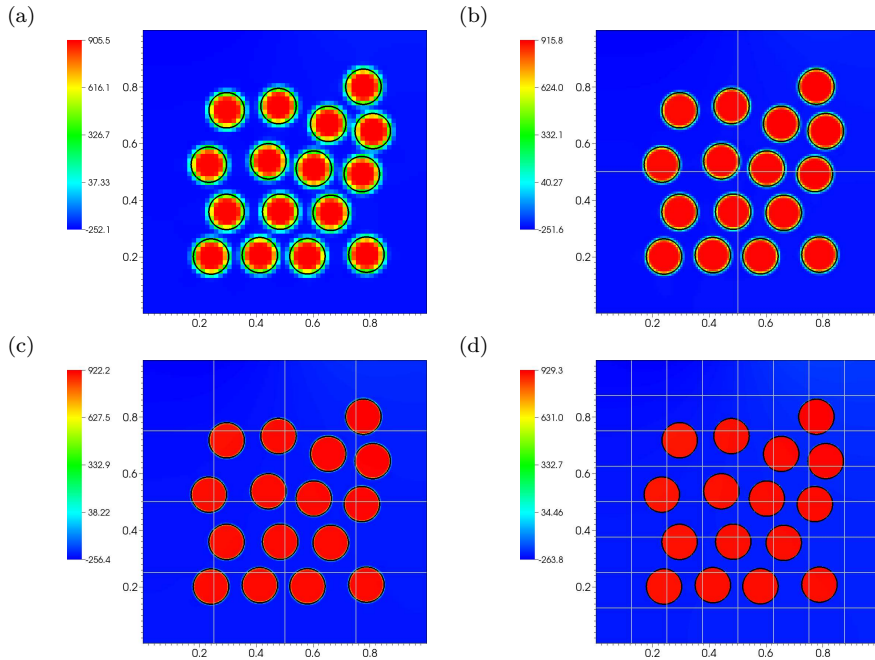


Fig. 12: Pressure field and parallel domain decomposition (indicated by gray boxes) for a suspension of 16 circular immersed interfaces with $\rho = 1$, $\mu = 1$, and $\gamma = 5$ for (a) $N = 64$, (b) $N = 128$, (c) $N = 256$, and (d) $N = 512$.

serial algorithm also yields good scalability. As also observed in the serial case, the Schur complement-based algorithm ultimately stalls for sufficiently small viscosities or sufficiently large elastic stiffnesses.

6 Discussion and conclusions

This paper has extended a geometric multigrid (GMG) preconditioning approach to semi-implicit formulations of the immersed boundary (IB) method [26, 27] in several important ways. First, we showed that the multiplicative “big-box” Vanka smoother previously developed by Guy et al. [27] can be recast as Richardson iterations preconditioned by a multiplicative Schwarz domain decomposition method, and we demonstrated that a restricted additive Schwarz (RAS) [13, 16] variant of this algorithm is also an effective smoother for the Stokes-IB systems that occur in this semi-implicit formulation. Although RAS yields convergence rates that are lower than multiplicative Schwarz, extending the smoother to a purely additive algorithm is crucial for deploying these methods in parallel computing environments because multiplicative domain decomposition methods impose a sequentiality that is not amenable to large-scale parallelization. Indeed, in the limit of large numbers of processors, it is

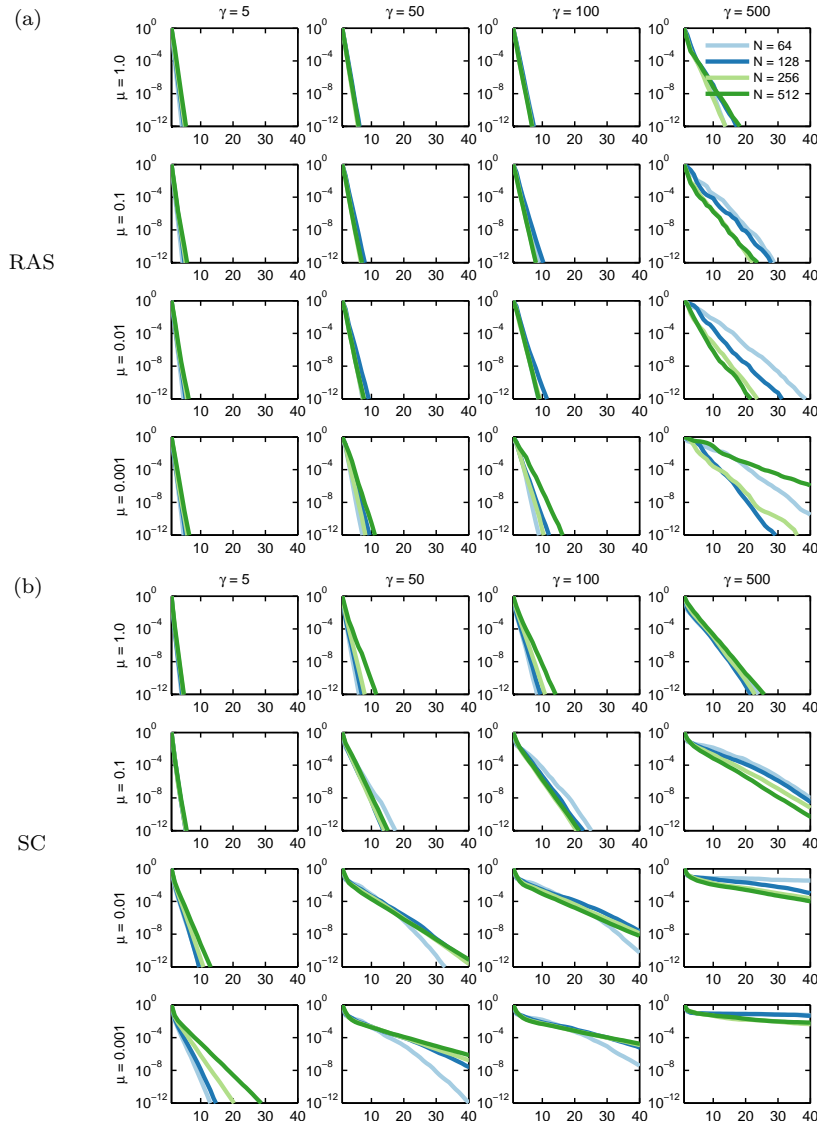


Fig. 13: Parallel scalability using the (a) RAS smoother and (b) Schur complement smoother at nonzero Reynolds number flow conditions for a suspension of elastic membranes (Sec. 5.3) using subdomains of size 16×16 with overlap widths of 4. We use $(N/64)^2$ processors for each case, so that the number of grid cells assigned to each processor remains fixed. Both solvers are essentially scalable for the lower Reynolds number cases, but as in the serial case, performance degrades with increasing elastic stiffness (γ) and decreasing viscosity (μ).

clear that an additive smoother that does not require or assume a particular order in which the subdomains is processed is essential to achieving good parallel scalability.

We further demonstrated that an even more effective smoother approach is obtained by considering an approximate block factorization of the Stokes-IB operator that appears in our semi-implicit formulation. What is remarkable about this Schur complement-based smoother is that it requires only a few iterations of point-relaxation smoothers on suitably constructed block operators for the velocity and pressure degrees of freedom. Consequently, the computational complexity of a single application of this smoother is comparable in complexity to optimal multigrid smoothers for much simpler systems such as isotropic Poisson problems. The SC smoother is also additive and well-suited for large-scale parallelization. Similar Stokes-type operator also appear in geodynamic applications that consider strong anisotropic viscosity variations (e.g. in the work of May et al. [35, 36, 19], which has proposed scalable multigrid preconditioners for such applications), and our Schur complement is similar to those used in this earlier work.

We performed extensive tests of the GMG algorithm using both RAS/RMS and SC smoothers. As in earlier work [27], we observe that solver performance degrades with increasing elastic stiffness. This study also reveals that the present solver approach degrades with increasing Reynolds number, with all methods ultimately failing for sufficiently small fluid viscosities. On the other hand, both the RAS/RMS and SC smoothers were shown to yield nearly optimal convergence rates at low Reynolds numbers and in Stokes flow conditions. Consequently, these methods may ultimately prove to offer practical solver strategies for important biological applications at the cellular and sub-cellular scales. The extension of this methodology to moderate-to-high Reynolds numbers remains important future work.

Although the present study considers only linear solver performance, this linear solver is implemented within a time stepping framework that supports both linear and nonlinear structural models. In the nonlinear case, we use a Newton-Krylov method [33], which requires solutions to systems of the form (6)–(9), but with \mathcal{F}_h replaced by a linearized force operator $\mathcal{K}_h = \left(\frac{\partial \mathcal{F}_h}{\partial \mathbf{X}}\right)\Big|_{\mathbf{X}=\mathbf{X}^{n+1}}$ for successive approximations to \mathbf{X}^{n+1} . We have found that because the configuration of the structure does not change very much within a time step, we generally can successfully use the configuration \mathbf{X}^n to construct a “lagged” preconditioner. JFNK generally appears to be quite effective so long as the underlying linear solver algorithm is effective. At present, however, the performance of both the linear and nonlinear implicit time stepping schemes lags that of our more mature explicit dynamics codes. In Stokes flow conditions and at very high stiffness, the implicit solver can yield wall-clock times comparable to our explicit solvers. In most other cases, however, the implicit solver generally remains a factor of 2–10 slower than the explicit solver, despite the fact that the implicit solver is able to use much larger time step sizes than the explicit solver. There are several reasons for this deficiency. First, the implicit code is substantially newer than our explicit code,

and although we have attempted to develop a reasonably well-optimized implementation, there is undoubtedly room for improvement (e.g. by switching to matrix-free operators where possible). Moreover, the RAS/RMS smoothers require the use of relatively large subdomains, which results in relatively large computational expenses. In some cases, the increased robustness of the implicit solvers may still justify their use, as with the implicit code, it is no longer necessary to carefully tune the time step size to avoid instabilities — a procedure that can substantially increase the time required to set up a complex model. We anticipate that further work, both in improving the algorithms and their implementations, will make the present scheme useful for low Reynolds number applications.

Treating moderate-to-high Reynolds numbers may require more than simply optimizing our implementation. One possibility would be to develop an alternative approximation to the Schur complement of Eq. (12) for use in the SC smoother. Alternatively, it may be necessary to reformulate the equations. For instance, one possibility is that instead of solving Eq. (11), we instead could solve

$$\begin{pmatrix} \mathcal{A} & \mathcal{G} & -\mathcal{S}_h^n \\ -\mathcal{D} & \mathbf{0} & \mathbf{0} \\ -\mathcal{J}_h^n & \mathbf{0} & \frac{1}{\Delta t}\mathcal{K}_h^{-1} \end{pmatrix} \begin{pmatrix} \mathbf{u}^{n+1} \\ p^{n+1} \\ \mathbf{F}^{n+1} \end{pmatrix} = \begin{pmatrix} \mathbf{g} \\ 0 \\ \mathbf{G} \end{pmatrix}. \quad (26)$$

A potential advantage of this formulation is that, for very large stiffnesses, the system is similar to a constrained formulation [31, 52]. Effective preconditioners have been developed for this class of problems [31, 52] and could potentially be extended to the case of FSI with stiff elastic structures. In the meantime, the development of effective, general-purpose preconditioners for implicit IB formulations with volumetric (codimension-0) structures remains an open problem.

A Spatial Discretization

This appendix briefly describes our spatial discretization of the IB equations (1)–(5), which is similar to that used in earlier studies [27, 22, 23].

A.1 Eulerian discretization

The Eulerian equations are approximated on a uniform Cartesian grid with grid spacing $h = \Delta x_1 = \Delta x_2$ using a staggered-grid discretization in which the Eulerian velocity \mathbf{u} and force \mathbf{f} are approximated at the centers of the Cartesian grid cell edges, and the Eulerian pressure is approximated at the centers of the grid cells. The cell centers are labeled using integer indices (i, j) , and the cell edges are labeled using shifted indices, i.e. $(i - \frac{1}{2}, j)$ for x_1 edges and $(i, j - \frac{1}{2})$ for x_2 edges. In this notation, $p_{i,j}$ indicates the approximation to $p(\mathbf{x}, t)$ at location $\mathbf{x}_{i,j}$, $(u_1)_{i-\frac{1}{2},j}$ indicates the approximation to the x_1 component of the velocity at location $\mathbf{x}_{i-\frac{1}{2},j}$, and $(f_2)_{i,j-\frac{1}{2}}$ indicates the approximation to the x_2 component of the force at location $\mathbf{x}_{i,j-\frac{1}{2}}$. Spatial Eulerian operators, including the scalar Laplacian ∇_h^2 and vector Laplacian ∇_h^2 , gradient ∇_h , and divergence $\nabla_h \cdot$ are discretized using standard second-order finite differences. Physical boundary conditions are treated in a manner described previously [21].

A.2 Lagrangian discretization

The Lagrangian force density $\mathbf{F}(\mathbf{s}, t)$ defined in Eq. (19) is discretized on a curvilinear mesh that is free to cut through the background Eulerian grid as the structure moves. The structure is discretized using a collection of Lagrangian nodes labeled by integer indices (l, m) , and we associate to each node curvilinear mesh spacings $(\Delta s_1, \Delta s_2)$. Simple finite difference approximations are used to evaluate the Lagrangian forces, as described previously [25, 24, 22, 23]. Specifically, an approximation to the derivative in the s_1 direction of a Lagrangian variable $\Phi(\mathbf{s}, t)$ is defined at a shifted ‘‘half-index’’ location by

$$(D_{s_1} \Phi)_{l+\frac{1}{2}, m} = \frac{\Phi_{l+1, m} - \Phi_{l, m}}{\Delta s_1}, \quad (27)$$

in which $\Phi_{l, m}$ approximates $\Phi(\mathbf{s}, t)$ at curvilinear mesh node $\mathbf{s}_{l, m}$. Our tests consider only fibers with a zero resting length, for which the fiber tension T and unit tangent vector $\boldsymbol{\tau}$ are also approximated at shifted locations by

$$T_{l+\frac{1}{2}, m} = \alpha \left\| (D_{s_1} \mathbf{X})_{l+\frac{1}{2}, m} \right\|, \quad (28)$$

$$\boldsymbol{\tau}_{l+\frac{1}{2}, m} = \frac{(D_{s_1} \mathbf{X})_{l+\frac{1}{2}, m}}{\left\| (D_{s_1} \mathbf{X})_{l+\frac{1}{2}, m} \right\|}. \quad (29)$$

Using these definitions of D_{s_1} , T , and $\boldsymbol{\tau}$, we compute an approximation to $F(\mathbf{s}_{l, m}, t)$ via

$$\mathbf{F}_{l, m} = (D_{s_1} (T \boldsymbol{\tau}))_{l, m}. \quad (30)$$

A.3 Lagrangian-Eulerian interaction

Interaction between Lagrangian and Eulerian variables is mediated by integral transforms (3) and (4). In the discrete version of the convolution equations, the singular Dirac delta kernel is replaced by a regularized kernel of the form $\delta_h(\mathbf{x}) = \prod_{i=1}^d \delta_h(x_i)$, in which the one-dimensional regularized kernel is $\delta_h(x_i) = \frac{1}{h} \phi\left(\frac{x_i}{h}\right)$. In this work, we use Peskin’s four-point regularized delta function [43], which is defined in terms of the basic kernel function

$$\phi(r) = \begin{cases} \frac{1}{8} \left(3 - 2|r| + \sqrt{1 + 4|r| - 4r^2} \right), & 0 \leq |r| < 1 \\ \frac{1}{8} \left(5 - 2|r| - \sqrt{-7 + 12|r| - 4r^2} \right), & 1 \leq |r| < 2 \\ 0, & 2 \leq |r|. \end{cases} \quad (31)$$

In two spatial dimensions, a discretized version of the force spreading equation (3) is used to obtain the Eulerian force density \mathbf{f} from $\mathbf{F} = (F_1, F_2)$ on the finest level of the locally refined Cartesian grid via

$$(f_1)_{i-\frac{1}{2}, j} = \sum_{l, m} (F_1)_{l, m} \delta_h(\mathbf{x}_{i-\frac{1}{2}, j} - \mathbf{X}_{l, m}) \Delta s_1 \Delta s_2, \quad (32)$$

$$(f_2)_{i, j-\frac{1}{2}} = \sum_{l, m} (F_2)_{l, m} \delta_h(\mathbf{x}_{i, j-\frac{1}{2}} - \mathbf{X}_{l, m}) \Delta s_1 \Delta s_2. \quad (33)$$

Similarly, the Eulerian fluid velocity $\mathbf{u} = (u_1, u_2)$ is interpolated to the curvilinear mesh on the finest grid level to obtain the structural velocity field $\mathbf{U} = (U_1, U_2)$ via

$$(U_1)_{l, m} = \sum_{i, j} (u_1)_{i-\frac{1}{2}, j} \delta_h(\mathbf{x}_{i-\frac{1}{2}, j} - \mathbf{X}_{l, m}) \Delta x_1 \Delta x_2, \quad (34)$$

$$(U_2)_{l, m} = \sum_{i, j} (u_2)_{i, j-\frac{1}{2}} \delta_h(\mathbf{x}_{i, j-\frac{1}{2}} - \mathbf{X}_{l, m}) \Delta x_1 \Delta x_2. \quad (35)$$

As in the continuous equations, we use the shorthand $\mathbf{f} = \mathcal{S}_h[\mathbf{X}]\mathbf{F}$ and $\mathbf{U} = \mathcal{J}_h[\mathbf{X}]\mathbf{u}$ for these discretized coupling operators. Moreover, so long as the operators are evaluated using the same structural configuration, $\mathcal{S}_h[\mathbf{X}] = \mathcal{J}_h^*[\mathbf{X}]$ because the same kernel function appears in both of the discretized integrals.

Acknowledgements

A.P.S.B. and R.D.G. gratefully acknowledge discussions with Gerry Puckett on related solvers used in geodynamics applications. We also gratefully acknowledge assistance from Barry Smith in profiling and optimizing the performance of the RAS and RMS smoothers used in this work.

Bibliography

References

1. IBAMR: An adaptive and distributed-memory parallel implementation of the immersed boundary method. <https://github.com/IBAMR/IBAMR>
2. Arnold, D.N., Boffi, D., Falk, R.S.: Quadrilateral $H(\text{div})$ finite elements. *SIAM J Numer Anal* **42**(6), 2429–2451 (2005)
3. Balay, S., Abhyankar, S., Adams, M.F., Brown, J., Brune, P., Buschelman, K., Dalcin, L., Eijkhout, V., Gropp, W.D., Kaushik, D., Knepley, M.G., McInnes, L.C., Rupp, K., Smith, B.F., Zampini, S., Zhang, H.: PETSc users manual. Tech. Rep. ANL-95/11 - Revision 3.6, Argonne National Laboratory (2015). URL <http://www.mcs.anl.gov/petsc>
4. Balay, S., Abhyankar, S., Adams, M.F., Brown, J., Brune, P., Buschelman, K., Dalcin, L., Eijkhout, V., Gropp, W.D., Kaushik, D., Knepley, M.G., McInnes, L.C., Rupp, K., Smith, B.F., Zampini, S., Zhang, H.: PETSc Web page. <http://www.mcs.anl.gov/petsc> (2015). URL <http://www.mcs.anl.gov/petsc>
5. Balay, S., Gropp, W.D., McInnes, L.C., Smith, B.F.: Efficient management of parallelism in object oriented numerical software libraries. In: E. Arge, A.M. Bruaset, H.P. Langtangen (eds.) *Modern Software Tools in Scientific Computing*, pp. 163–202. Birkhäuser Press (1997)
6. Bhalla, A.P.S., Bale, R., Griffith, B.E., Patankar, N.A.: A unified mathematical framework and an adaptive numerical method for fluid-structure interaction with rigid, deforming, and elastic bodies. *J Comput Phys* **250**, 446–476 (2013)
7. Bhalla, A.P.S., Bale, R., Griffith, B.E., Patankar, N.A.: Fully resolved immersed electrohydrodynamics for particle motion, electrolocation, and self-propulsion. *J Comput Phys* **256**, 88–108 (2014)
8. Bhalla, A.P.S., Griffith, B.E., Patankar, N.A., Donev, A.: A minimally-resolved immersed boundary model for reaction-diffusion problems. *J Chem Phys* **139**(21), 214,112 (15 pages) (2013)
9. Boffi, D., Gastaldi, L., Heltai, L., Peskin, C.S.: On the hyper-elastic formulation of the immersed boundary method. *Comput Meth Appl Mech Eng* **197**(25–28), 2210–2231 (2008)
10. Borazjani, I., Ge, L., Sotiropoulos, F.: Curvilinear immersed boundary method for simulating fluid structure interaction with complex 3D rigid bodies. *J Comput Phys* **227**(16), 7587–7620 (2008)
11. Briggs, W.L., Henson, V.E., McCormick, S.F.: *A Multigrid Tutorial*, Second Edition. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA (2000)
12. Cai, M., Nonaka, A., Bell, J.B., Griffith, B.E., Donev, A.: Efficient variable-coefficient finite-volume Stokes solvers. *Comm Comput Phys* **16**(5), 1263–1297 (2014)
13. Cai, X.C., Sarkis, M.: A restricted additive schwarz preconditioner for general sparse linear systems. *SIAM J Sci Comput* **21**(2), 792–797 (1999)

14. Ceniceros, H.D., Fisher, J.E., Roma, A.M.: Efficient solutions to robust, semi-implicit discretizations of the immersed boundary method. *J Comput Phys* **228**(19), 7137–7158 (2009)
15. Colella, P., Woodward, P.R.: The piecewise parabolic method (PPM) for gas-dynamical simulations. *J Comput Phys* **54**(1), 174–201 (1984)
16. Efstathiou, E., Gander, M.J.: Why restricted additive schwarz converges faster than additive schwarz. *BIT* **43**(5), 945–959 (2003)
17. Elman, H., Howle, V.E., Shadid, J., Shuttleworth, R., Tuminaro, R.: A taxonomy and comparison of parallel block multi-level preconditioners for the incompressible Navier-Stokes equations. *J Comput Phys* **227**(3), 1790–1808 (2008)
18. Flamini, V., DeAnda, A., Griffith, B.E.: Immersed boundary-finite element model of fluid-structure interaction in the aortic root. *Theor Comput Fluid Dynam* **30**(1), 139–164 (2016)
19. Furuichi, M., May, D.A., Tackley, P.J.: Development of a stokes flow solver robust to large viscosity jumps using a schur complement approach with mixed precision arithmetic. *J Comput Phys* **230**(24), 8835–8851 (2011)
20. Glowinski, R., Pan, T.W., Hesla, T.I., Joseph, D.D.: A distributed lagrange multiplier/fictitious domain method for particulate flows. *Int J Multiphase Flow* **25**(5), 755–794 (1999)
21. Griffith, B.E.: An accurate and efficient method for the incompressible Navier-Stokes equations using the projection method as a preconditioner. *J Comput Phys* **228**(20), 7565–7595 (2009)
22. Griffith, B.E.: Immersed boundary model of aortic heart valve dynamics with physiological driving and loading conditions. *Int J Numer Meth Biomed Eng* **28**(3), 317–345 (2012)
23. Griffith, B.E.: On the volume conservation of the immersed boundary method. *Commun Comput Phys* **12**(2), 401–432 (2012)
24. Griffith, B.E., Hornung, R.D., McQueen, D.M., Peskin, C.S.: An adaptive, formally second order accurate version of the immersed boundary method. *J Comput Phys* **223**(1), 10–49 (2007)
25. Griffith, B.E., Peskin, C.S.: On the order of accuracy of the immersed boundary method: Higher order convergence rates for sufficiently smooth problems. *J Comput Phys* **208**(1), 75–105 (2005)
26. Guy, R.D., Philip, B.: A multigrid method for a model of the implicit immersed boundary equations. *Comm Comput Phys* **12**(2), 378–400 (2012)
27. Guy, R.D., Phillip, B., Griffith, B.E.: Geometric multigrid for an implicit-time immersed boundary method. *Adv Comput Math* **41**(3), 635–662 (2015)
28. Hornung, R.D., Kohn, S.R.: Managing application complexity in the SAMRAI object-oriented framework. *Concurrency Comput Pract Ex* **14**(5), 347–368 (2002)
29. Hou, T.Y., Shi, Z.: An efficient semi-implicit immersed boundary method for the Navier-Stokes equations. *J Comput Phys* **227**(10), 8968–8991 (2008)
30. Hou, T.Y., Shi, Z.: Removing the stiffness of elastic force from the immersed boundary method for the 2D Stokes equations. *J Comput Phys* **227**(21), 9138–9169 (2008)
31. Kallemov, B., Bhalla, A.P.S., Griffith, B.E., Donev, A.: An immersed boundary method for rigid bodies. *Comm Appl Math Comput Sci* **11**(1), 79–141 (2016)
32. Kim, Y., Peskin, C.S.: Penalty immersed boundary method for an elastic boundary with mass. *Phys Fluid* **19**, 053,103 (18 pages) (2007)
33. Knoll, D.A., Keyes, D.E.: Jacobian-free Newton-Krylov methods: A survey of approaches and applications. *J Comput Phys* **193**(2), 357–397 (2004)
34. Lai, M.C., Peskin, C.S.: An immersed boundary method with formal second-order accuracy and reduced numerical viscosity. *J Comput Phys* **160**(2), 705–719 (2000)
35. May, D.A., Brown, J., Le Pourhiet, L.: A scalable, matrix-free multigrid preconditioner for finite element discretizations of heterogeneous stokes flow. *Comput Meth Appl Mech Eng* **290**, 496–523 (2015)
36. May, D.A., Moresi, L.: Preconditioned iterative methods for stokes flow problems arising in computational geodynamics. *Physics of the Earth and Planetary Interiors* **171**(1), 33–47 (2008)

37. Mayo, A.A., Peskin, C.S.: An implicit numerical method for fluid dynamics problems with immersed elastic boundaries. In: A.Y. Cheer, C.P. van Dam (eds.) *Fluid Dynamics in Biology: Proceedings of an AMS-IMS-SIAM Joint Summer Research Conference, Contemporary Mathematics*, vol. 140, pp. 261–277. American Mathematical Society, Providence, RI, USA (1993)
38. Mittal, R., Dong, H., Bozkurtas, M., Najjar, F., Vargas, A., von Loebbecke, A.: A versatile sharp interface immersed boundary method for incompressible flows with complex boundaries. *J Comput Phys* **227**(10), 4825–4852 (2008)
39. Mori, Y., Peskin, C.S.: Implicit second order immersed boundary methods with boundary mass. *Comput Meth Appl Mech Eng* **197**(25–28), 2049–2067 (2008)
40. Newren, E.P., Fogelson, A.L., Guy, R.D., Kirby, R.M.: Unconditionally stable discretizations of the immersed boundary equations. *J Comput Phys* **222**(2), 702–719 (2007)
41. Peskin, C.S.: Flow patterns around heart valves: a numerical method. *J Comput Phys* **10**(2), 252–271 (1972)
42. Peskin, C.S.: Numerical analysis of blood flow in the heart. *J Comput Phys* **25**(3), 220–252 (1977)
43. Peskin, C.S.: The immersed boundary method. *Acta Numer* **11**, 479–517 (2002)
44. Rider, W.J., Greenough, J.A., Kamm, J.R.: Accurate monotonicity- and extrema-preserving methods through adaptive nonlinear hybridizations. *J Comput Phys* **225**(2), 1827–1848 (2007)
45. Saad, Y.: A flexible inner-outer preconditioned GMRES algorithm. *SIAM J Sci Comput* **14**(2), 461–469 (1993)
46. SAMRAI: Structured Adaptive Mesh Refinement Application Infrastructure. <http://www.llnl.gov/CASC/SAMRAI>
47. Smith, B.F., Bjørstad, P.E., Gropp, W.D.: *Domain Decomposition: Parallel Multilevel Methods for Elliptic Partial Differential Equations*. Cambridge University Press (1996)
48. Stein, D.B., Guy, R.D., Thomases, B.: Immersed boundary smooth extension: A high-order method for solving pde on arbitrary smooth domains using fourier spectral methods. *J Comput Phys* **304**, 252–274 (2016)
49. Trottenberg, U., Oosterlee, C., Schüller, A.: *Multigrid*. Elsevier, San Diego, CA, USA (2001)
50. Tseng, Y.H., Ferziger, J.H.: A ghost-cell immersed boundary method for flow in complex geometry. *J Comput Phys* **192**(2), 593–623 (2003)
51. Uhlmann, M.: An immersed boundary method with direct forcing for the simulation of particulate flows. *J Comput Phys* **209**(2), 448–476 (2005)
52. Usabiaga, F.B., Kallemov, B., Delmotte, B., Bhalla, A.P.S., Griffith, B.E., Donev, A.: Hydrodynamics of suspensions of passive and active rigid particles: A rigid multiblob approach. To appear in *Comm Appl Math Comput Sci*
53. Vanka, S.P.: Block-implicit multigrid solution of Navier-Stokes equations in primitive variables. *J Comput Phys* **65**(1), 138–158 (1986)
54. Zhang, L., Gerstenberger, A., Wang, X., Liu, W.K.: Immersed finite element method. *Comput Meth Appl Mech Eng* **193**(21–22), 2051–2067 (2004)
55. Zhang, Q., Guy, R.D., Philip, B.: A projection preconditioner for solving the implicit immersed boundary equations. *Numer Math Theor Meth Appl* **7**(4), 473–498 (2014)