Computers and Fluids 214 (2021) 104764

Contents lists available at ScienceDirect

Computers and Fluids

journal homepage: www.elsevier.com/locate/compfluid

An octree-based immersogeometric approach for modeling inertial migration of particles in channels

Songzhe Xu[a], Boshun Gao[a], Alec Lofquist[a], Milinda Fernando[b], Ming-Chen Hsu[a], Hari Sundar[b,*], Baskar Ganapathysubramanian[a,*]

[a] Department of Mechanical Engineering, Iowa State University, 2080 Black Engineering, Ames, IA 50011, USA
[b] School of Computing, University of Utah, 50 Central Campus Dr, Salt Lake City, UT 84112, USA

ARTICLE INFO

Article history:
Received 31 July 2020
Accepted 11 October 2020
Available online 21 October 2020

Keywords:
Octree-based mesh
Adaptive mesh refinement
Immersogeometric analysis
Particle inertial migration

ABSTRACT

In this paper, we develop a scalable, adaptively refined, octree-based finite element approach with immersogeometric analysis to track inertial migration of particles in microchannels. Fluid physics is modeled using a residual-based variational multiscale method, and the particle movement is modeled as rigid body motion. A parallel, hierarchically refined octree mesh is employed as the background mesh, on which a variationally consistent immersogeometric formulation is adopted for tracking the particle motion in the fluid. Adaptations of immersogeometric analysis on an octree background mesh are developed to enable efficient searching of background element for a given surface quadrature point, as well as a distribution of surface quadrature points over processors to reduce memory overhead and better parallelize the surface assembly. An octree-based adaptive mesh refinement algorithm adapted to in-out test in the immersogeometric approach is also developed. The validation of our octree-based immersogeometric approach is carried out using a benchmark case of a sphere settling in quiescent fluid, with good agreement presented. In addition, good strong (and weak) scalability on supercomputing resources for this benchmark case up to 16,384 processes is demonstrated. The proposed method is further deployed for exploring particle migration in straight and converging-diverging channels. This example illustrates the potential of the octree-based immersogeometric approach for efficiently tracking particle motion in complex channel flows – a problem with a diverse array of applications.

1. Introduction

Control and localization of particles (e.g. cells and precipitates) in flow is useful in biological processing, chemical reaction control, and for creating structured materials [1,2]. One promising approach to control localization (or 'focusing') of particles of different sizes is via sequential placement of obstacles in microchannels. This is based on the idea that obstacles produce different forces on particles with distinct sizes when inertial flow conditions are considered (i.e., when the Reynolds number based on the channel hydraulic diameter is greater than 5) [3]. However, highly accurate force calculations are required to design devices that can exploit these small variations in forces, which essentially becomes a comput55ational exercise. The availability of validated approaches to compute these dynamics can enable diverse applications in separation, concentration, and sorting of cells and biomolecules with high specificity.

The general problem of force (and trajectory) calculation on a moving particle in channel flow can be framed as a canonical problem. The canonical problem is to track the lateral migration of a single, rigid particle as it traverses a microchannel that is decorated with a pillar obstacle (see Fig. 1). We consider a rigid particle (of size a) moving in an incompressible fluid inside a channel. We are interested in tracking the motion and lateral forces acting on this particle. This is a challenging computational problem due to the full fluid–solid coupling, associated small time step sizes, and adaptive (re)meshing requirements. The full fluid–solid coupling is needed due to the finite Reynolds number ($5 \leq Re \leq 100$) which necessitates solving the full Navier–Stokes equations. Furthermore, the construction of migration maps for particles (i.e. how does a particle traverse a decorated channel, which is a function of initial particle release locations, particle sizes and flow rates) requires several hundreds of simulations tracking individual particles under various configurations (of release locations and particle sizes).

* Corresponding authors.
E-mail addresses: hari@cs.utah.edu (H. Sundar), baskarg@iastate.edu (B. Ganapathysubramanian).

https://doi.org/10.1016/j.compfluid.2020.104764
0045-7930/

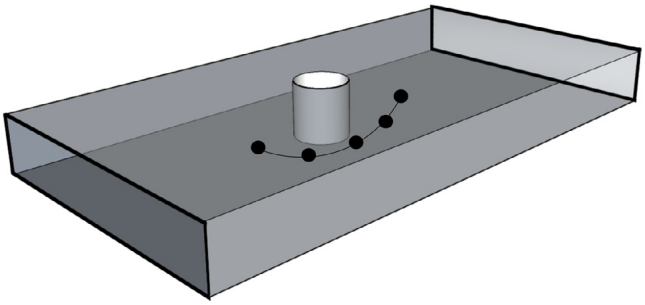

Fig. 1. Canonical problem: A rigid particle traversing a microchannel decorated with a pillar obstacle under inertial flow conditions.

To track particles in decorated channels, immersed boundary methods [4,5] are a popular class of approaches, since these methods exhibit greater geometric flexibility than their boundary-fitted counterparts, and simplify (re)meshing process especially when the object is moving. The immersed boundary method embeds the solid geometry into a background mesh without the need for conforming the solid and fluid meshes. As a result, it becomes computationally convenient to track the motion of arbitrary particles while avoiding a cumbersome boundary-fitted (re)meshing process. In the context of finite elements, several adaptations of immersed methods have been explored for the simulation of fluid interacting with moving objects [6–13]. Among these immersed methods, immersogeometric analysis [14–18] (IMGA) is a promising approach for accurately predicting flow results by faithfully capturing the immersed geometry using adaptively refined quadrature rules in the intersected elements, and weakly enforcing the Dirichlet boundary conditions on the surface of the immersed object using an extension of Nitsche's method [19]. Our prior work has shown that the IMGA is a viable approach to track particles in microchannels [20], and we extend the IMGA approach to 3D simulations in this work.

The IMGA approach incorporates a residual-based variational multiscale (VMS) formulation [21] to model fluid physics, which has been a successful approach to model complex flow phenomena. The VMS approach is analogous to a large eddy simulation (LES) model which uses variational projections in place of the traditional filtered equations in LES and focuses on modeling the fine-scale equations. The VMS approach does not employ any eddy viscosity, and has been successfully used to perform accurate flow condition agnostic (laminar or turbulent) simulations. It has been extended to a wide range of engineering applications, such as buoyancy driven flows, particle laden flows, fluid–structure interaction, multiphase and free-surface flows, space-time thermal flows, magnetohydrodynamic flows, and compressible flows [22–30]. While the IMGA is a promising approach to model particle migration in channels, it is challenging to optimize the computational efficiency on an unstructured background mesh. Thus, in this work, we proposed a new computational framework that performs large-scale finite element computations using IMGA on an octree-based background mesh due to its convenience and efficiency of fast adaptive mesh refinement and partitioning compared with traditional unstructured meshing approaches[1]

Octree-based meshing has been successfully applied to finite element computations for many engineering problems [33–38]. Specifically in this work, we employ the optimized parallel octree-based meshing library, Dendro, which has been deployed for simulating binary black hole inspiral, solving PDEs in 4D space-time octrees, reconstructing the motion of the ventricular walls in cardiac images, and computing information theoretic similarity measures for medical images [39–43]. Some of the features of Dendro include a bottom-up construction of octrees and a 2:1 balancing on distributed architectures [44,45], and a space filling curve partitioning for load balancing [46–48]. While the concept of octree-based adaptive space partitions is well studied, developing such methods for the IMGA on large distributed systems is novel. In addition, adaptations of existing IMGA on the octree background mesh are developed to enable an efficient searching of background element for a given surface quadrature point, as well as a distribution of surface quadrature points over processors to reduce memory overhead and better parallelize the surface assembly in the IMGA. The original octree-based adaptive mesh refinement algorithm is also adjusted for the in-out test in the IMGA. Finally, the proposed octree-based IMGA framework is applied to the simulations of tracking particles in channels to demonstrate its accuracy and scalability.

The paper is organized as follows. In Section 2, we summarize the formulations of the incompressible Navier–Stokes equations and particle motion. Section 3 describes the semi-discrete formulations and time stepping. Section 4 briefs the octree-based mesh implementations and adaptations of extending IMGA on the octree-based mesh. In Section 5, we validate the framework and show scaling results using a case of sphere dropping in a quiescent fluid. We also show a few applications of this framework for tracking particles in different microchannels. Finally, we draw conclusions and motivate future research in Section 6.

2. Governing equations for particles moving in fluids

We consider the non-dimensional two-way coupled governing equations describing the interactions between the particle and the fluid in the channel.

2.1. Incompressible Navier–Stokes equations

The flow is described by the dimensionless Navier–Stokes equations posed on a fluid domain Ω_t:

$$\left(\frac{\partial \mathbf{u}}{\partial t} + \mathbf{u}\cdot\nabla\mathbf{u}\right) - \nabla\cdot\boldsymbol{\sigma} = \mathbf{0}, \tag{1}$$

$$\nabla\cdot\mathbf{u} = 0, \tag{2}$$

where $t$is the time, and $\mathbf{u}$is the flow velocity. The stress and strain-rate tensors are defined respectively as

$$\boldsymbol{\sigma}(\mathbf{u}, p) = -p\mathbf{I} + 2\frac{1}{Re}\boldsymbol{\varepsilon}(\mathbf{u}), \tag{3}$$

$$\boldsymbol{\varepsilon}(\mathbf{u}) = \frac{1}{2}\left(\nabla\mathbf{u} + \nabla\mathbf{u}^{\mathrm{T}}\right), \tag{4}$$

where $p$is the pressure, $\mathbf{I}$is an identity tensor, and $Re$is the channel hydraulic diameter based Reynolds number. The problem (1)–(4) is accompanied by suitable boundary conditions, defined on the boundary of the fluid domain, $\Gamma_t = \Gamma_t^{\mathrm{D}} \cup \Gamma_t^{\mathrm{N}}$:

$$\mathbf{u} = \mathbf{u_g} \qquad \text{on } \Gamma_t^{\mathrm{D}}, \tag{5}$$

$$-p\,\mathbf{n} + 2\frac{1}{Re}\boldsymbol{\varepsilon}(\mathbf{u})\,\mathbf{n} = \mathbf{h} \qquad \text{on } \Gamma_t^{\mathrm{N}}, \tag{6}$$

where $\mathbf{u_g}$denotes the prescribed velocity at the Dirichlet boundary Γ_t^{D},$\mathbf{h}$is the traction vector at the Neumann boundary Γ_t^{N},and $\mathbf{n}$is the unit normal vector pointing in the wall-outward direction.

[1] This work is based on the thesis work of the authors, Xu [31] and Lofquist [32].

2.2. Particle motion

The particle is modeled as a rigid body. We denote the dimensionless kinematic state of the object as $\mathbf{Y}$,and the equations of motion may be written in a Lagrange frame of reference as follows [49]:

$$\mathbf{Y} = \begin{bmatrix} \mathbf{x}^c \\ \mathbf{R} \\ \mathbf{P} \\ \mathbf{L} \end{bmatrix}, \qquad \dot{\mathbf{Y}} = \frac{d\mathbf{Y}}{dt} = \begin{bmatrix} \mathbf{v}^c \\ \boldsymbol{\omega}^* \mathbf{R} \\ \mathbf{F} \\ \mathbf{T} \end{bmatrix}, \tag{7}$$

where

$$\mathbf{v}^c = \frac{\mathbf{P}}{M}, \qquad \boldsymbol{\omega} = \mathbf{I}^{-1}\mathbf{L}, \tag{8}$$

and

$$\mathbf{I} = \mathbf{R}\mathbf{I}_{\mathbf{ini}}\mathbf{R}^T, \quad \boldsymbol{\omega}^* = \begin{pmatrix} 0 & -\omega_z & \omega_y \\ \omega_z & 0 & \omega_x \\ -\omega_y & \omega_x & 0 \end{pmatrix}. \tag{9}$$

In Eqs. (7)–(9), $\mathbf{x}^c$is the position of the center of mass of the object, $\mathbf{v}^c$is the velocity of the center of mass of the object, $\mathbf{R}$is the rotation matrix mapping from initial configuration to current configuration, $\boldsymbol{\omega}$is the object's angular velocity, $\mathbf{P}$and $\mathbf{L}$are linear and angular momentum, respectively. $M$and $\mathbf{I}$are the mass and inertia tensor of the object, both of which are non-dimensinalized using the fluid density and characteristic length scales. and the inertia tensor $\mathbf{I}$can be further defined using $\mathbf{I}_{\mathbf{ini}}$in initial configuration, which is constant during the motion. In Eq. (7), $\mathbf{F}$and $\mathbf{T}$are the integral of force and torque acting on the particle surface which are computed from the solution of the fluid field, and defined as follows:

$$\mathbf{F} = \oint_{\Gamma_t^I} \boldsymbol{\sigma}(\mathbf{u}, p)\mathbf{n} d\Gamma_t, \qquad \mathbf{T} = \oint_{\Gamma_t^I} \mathbf{r} \times (\boldsymbol{\sigma}(\mathbf{u}, p)\mathbf{n}) d\Gamma_t, \tag{10}$$

where $\mathbf{r}$is the distance vector from the particle's center of mass to its surface points in the current configuration, $\Gamma_t^I$is the particle boundary, and the coordinates $\mathbf{x}$and velocities $\mathbf{v}$of its surface points are computed as

$$\mathbf{x} = \mathbf{x}^c + \mathbf{r}, \qquad \mathbf{v} = \mathbf{v}^c + \boldsymbol{\omega} \times \mathbf{r}. \tag{11}$$

Finally, $\mathbf{n}$is the unit normal vector that points outward from the particle surface. Note, most microfluidic applications involve neutrally buoyant particles (i.e. the density of particle matches the density of fluid). Density matching allows us to omit the buoyancy term in Eq. (10). It is trivial to include buoyancy for case of unmatched densities.

3. Semi-discrete formulation and time discretization

3.1. Variational multiscale formulation

Consider a collection of disjoint elements $\{\Omega_t^e\}$,$\cup_e \Omega_t^e \subset \mathbb{R}^d$. The fluid domain is covered by the closure of the collection: $\Omega_t \subset \cup_e \overline{\Omega_t^e}$. Note that $\Omega_t^e$is not necessarily a subset of $\Omega_t$with the immersed boundary method. Let $\mathcal{V}_u^h$and $\mathcal{V}_p^h$be the finite-dimensional spaces of discrete test functions and trial solutions for velocity and pressure, which are denoted as superscript h,and represent resolved scales (coarse scale) produced by the finite element discretization. The strong problem (1)–(6) may be recast in a weak form and posed over these discrete spaces to produce the following semi-discrete problem using the VMS modeling approach: Find $\mathbf{u}^h \in \mathcal{V}_u^h$and $p^h \in \mathcal{V}_p^h$such that for all $\mathbf{w}^h \in \mathcal{V}_u^h$and $q^h \in \mathcal{V}_p^h$:

$$B^{\mathrm{VMS}}\left(\{\mathbf{w}^h, q^h\}, \{\mathbf{u}^h, p^h\}\right) - F^{\mathrm{VMS}}\left(\{\mathbf{w}^h, q^h\}\right) = 0. \tag{12}$$

The bilinear form $B^{\mathrm{VMS}}$and the load vector $F^{\mathrm{VMS}}$are given as

$$\begin{aligned}
&B^{\mathrm{VMS}}\left(\{\mathbf{w}^h, q^h\}, \{\mathbf{u}^h, p^h\}\right) \\
&= \int_{\Omega_t} \mathbf{w}^h \cdot \left(\frac{\partial \mathbf{u}^h}{\partial t} + \mathbf{u}^h \cdot \nabla \mathbf{u}^h\right) d\Omega_t \\
&+ \int_{\Omega_t} \nabla \mathbf{w}^h : \boldsymbol{\sigma}\left(\mathbf{u}^h, p^h\right) d\Omega_t \\
&+ \int_{\Omega_t} q^h \nabla \cdot \mathbf{u}^h \, d\Omega_t \\
&- \sum_e \int_{\Omega_t^e \cap \Omega_t} \left(\mathbf{u}^h \cdot \nabla \mathbf{w}^h + \nabla q^h\right) \cdot \mathbf{u}' \, d\Omega_t \\
&- \sum_e \int_{\Omega_t^e \cap \Omega_t} p' \nabla \cdot \mathbf{w}^h \, d\Omega_t \\
&+ \sum_e \int_{\Omega_t^e \cap \Omega_t} \mathbf{w}^h \cdot (\mathbf{u}' \cdot \nabla \mathbf{u}^h) \, d\Omega_t \\
&- \sum_e \int_{\Omega_t^e \cap \Omega_t} \nabla \mathbf{w}^h : \left(\mathbf{u}' \otimes \mathbf{u}'\right) d\Omega_t,
\end{aligned} \tag{13}$$

and

$$F^{\mathrm{VMS}}\left(\{\mathbf{w}^h, q^h\}\right) = \int_{\Omega_t} \mathbf{w}^h \cdot \mathbf{f} \, d\Omega_t + \int_{\Gamma_t^N} \mathbf{w}^h \cdot \mathbf{h} \, d\Gamma_t, \tag{14}$$

where the variables with superscript primes denote the unsolved scales (fine scale) that need to be modeled, and their effect is added onto the coarse scale. $\mathbf{u}'$is defined as

$$\mathbf{u}' = -\tau_{\mathrm{M}}\left(\frac{\partial \mathbf{u}^h}{\partial t} + \mathbf{u}^h \cdot \nabla \mathbf{u}^h - \mathbf{f} - \nabla \cdot \boldsymbol{\sigma}\left(\mathbf{u}^h, p^h\right)\right), \tag{15}$$

and $p'$is given by

$$p' = -\tau_{\mathrm{C}} \nabla \cdot \mathbf{u}^h. \tag{16}$$

Here, $\mathbf{u}'$and $p'$are approximated by the residuals of momentum equation and continuity equation, respectively, and $\tau_M$and $\tau_C$are corresponding coefficients with the definitions in Bazilevs et al. [21]. Eqs. (12)–(16) feature the VMS formulation of Navier–Stokes equations of incompressible flows. The additional terms added onto the standard weak Galerkin form can be interpreted as a combination of streamline/upwind Petrov–Galerkin (SUPG) stabilization and VMS large-eddy simulation of turbulence modeling.

3.2. Immersogeometric analysis

The no-slip boundary condition (which is a Dirichlet boundary condition) on an immersed particle surface is converted into an equivalent Neumann condition in the sense of the Nitsche's method [50]. We perform a surface integral over the immersed boundary to weakly impose the Dirichlet boundary condition [19,51,52]. Assuming the immersed boundary $\Gamma_t^I$is decomposed into N_{eb}surface elements each denoted by $(\Gamma_t^I)_b$,the semi-discrete problem becomes

$$\begin{aligned}
&B^{\mathrm{VMS}}\left(\{\mathbf{w}^h, q^h\}, \{\mathbf{u}^h, p^h\}\right) - F^{\mathrm{VMS}}\left(\{\mathbf{w}^h, q^h\}\right) \\
&- \sum_{b=1}^{N_{eb}} \int_{(\Gamma_t^I)_b} \mathbf{w}^h \cdot \left(-p^h \mathbf{n} + 2\frac{1}{Re}\boldsymbol{\varepsilon}(\mathbf{u}^h)\mathbf{n}\right) d\Gamma \\
&- \sum_{b=1}^{N_{eb}} \int_{(\Gamma_t^I)_b} \left(2\frac{1}{Re}\boldsymbol{\varepsilon}(\mathbf{w}^h)\mathbf{n} + q^h \mathbf{n}\right) \cdot \left(\mathbf{u}^h - \mathbf{v}\right) d\Gamma \\
&+ \sum_{b=1}^{N_{eb}} \int_{(\Gamma_t^I)_b} \tau^B \mathbf{w}^h \cdot \left(\mathbf{u}^h - \mathbf{v}\right) d\Gamma \;=\; 0,
\end{aligned} \tag{17}$$

where $\mathbf{n}$is the normal vector of the immersed boundary. The boundary terms added to the governing equations are the second, third and last lines in Eq. (17), and a detailed interpretation of different terms can be found in Bazilevs and Hughes [19].

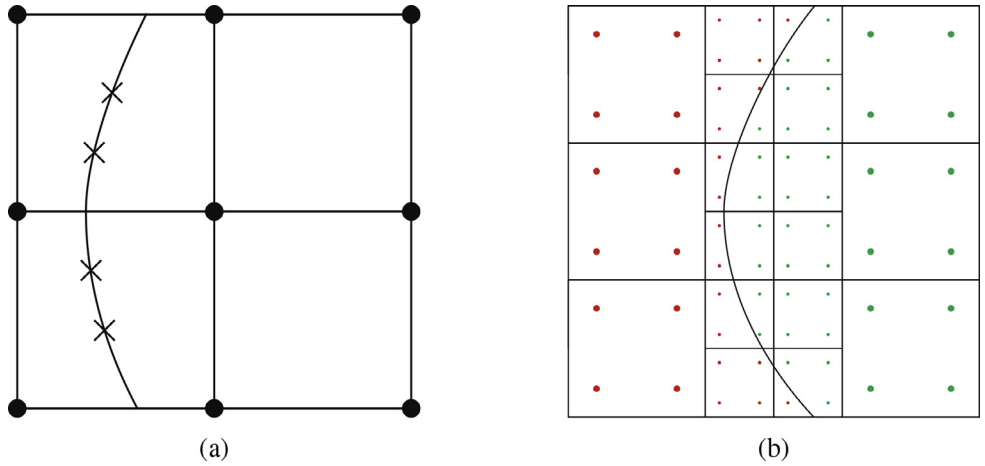

Fig. 2. Implementation of the immersogeometric method. (a) A schematic (2D example) showing how the surface assembly of IMGA is performed. The surface mesh is used to identify surface Gauss points (the 'X' locations). The immersed boundary condition terms (i.e. the last three terms in Eq. (17)) are computed at these surface Gauss points, and then distributed to their background nodes. (b) A schematic (2D example) of the volume assembly in the IMGA method. An in-out test is performed to identify whether each volume Gauss point lies inside the particle (green points) or inside the fluid (red points). Only the Gauss points in the fluid domain are used to assemble the elemental matrices. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

Only the penalty-like stabilization parameter, τ^B,is a heuristic that has to be appropriately chosen. We use the definition proposed in Wu et al. [53], which scales the stabilization parameter as $\tau^B = \max\{C^B_{\text{inv}} h/\Delta t, C^B_{\text{vis}}/(Re\, h)\}$,where the C^B's are positive dimensionless constants, $h$is the size of the cut element, and $\Delta t$is the time step size.

The boundary terms are imposed onto the surface Gauss points, which are then interpolated by their background Cartesian grids as shown in Fig. 2(a). In this way we can apply the Dirichlet boundary condition on the immersed boundary of the object. The implementation of the IMGA requires some refinement of the background mesh across the immersed surface to better capture the shape of the interface. The volume assembly of IMGA is accomplished by using selective quadrature (i.e. only using the Gauss points that lie in the fluid and not inside the immersed particle). This necessitates performing an in-out test to determine the Gauss points inside the fluid domain (red points) on which we assemble, while discarding the Gauss points inside the object (green points), as shown in Fig. 2(b). Note that the in-out test is also required for the adaptive refinement, and selective quadrature is performed only for the cut elements during volume assembly.

The particle evolution is then computed by evaluating the force and torque exerted from the fluids on the particle. Finally, when the particle is moving, freshly-cleared nodes, i.e., some background mesh nodes that are inside the object at one time step, but are in the fluid domain at the next time step will occur. These mesh nodes have no fluid history and require interpolation from their neighbors. We refer readers to our previous work [20] for a detailed treatment.

3.3. Time stepping for fluids and particle motion

The time-dependent Navier–Stokes equations are solved using a backward Euler implicit scheme as follows

$$\frac{\partial \mathbf{u}}{\partial t} = \frac{\mathbf{u}^n - \mathbf{u}^{n-1}}{\Delta t} = \mathcal{L}(\mathbf{u}^n, p^n), \tag{18}$$

where the operator $\mathcal{L}(\mathbf{u}^n, p^n)$represents all the other terms except the time-dependent term evaluated at the current time step in Eq. (1). $\Delta t$is chosen to respect the CFL condition[2]. The (non)linear solution procedure is taken care by PETSc [54]. We utilize the SNES construct (line search quasi-Newton), which uses the KSP construct, specifically the stabilized bi-conjugate gradient (BCGS) solver. An additive Schwarz preconditioner (ASM) is used to enable parallel preconditioning and solving on decomposed sub-domains.

An explicit forward Euler time-stepper is used to update the particle location and velocity. In the discrete form we have

$$\mathbf{Y}^{n+1} = \mathbf{Y}^n + \Delta t \dot{\mathbf{Y}}^n. \tag{19}$$

$\mathbf{F}^n$and $\mathbf{T}^n$at each time step are discretized in space and computed with weakly imposed boundary conditions as follows

$$\mathbf{F}^n = \sum_{b=1}^{N_{eb}} \int_{(\Gamma_t^1)_b} \boldsymbol{\sigma}(\mathbf{u}^n, p^n)\mathbf{n} d\Gamma - \sum_{b=1}^{N_{eb}} \int_{(\Gamma_t^1)_b} \tau^B(\mathbf{u}^n - \mathbf{v}^n) d\Gamma, \tag{20}$$

$$\mathbf{T}^n = \sum_{b=1}^{N_{eb}} \int_{(\Gamma_t^1)_b} \mathbf{r}\times(\boldsymbol{\sigma}(\mathbf{u}^n, p^n)\mathbf{n}) d\Gamma - \sum_{b=1}^{N_{eb}} \int_{(\Gamma_t^1)_b} \mathbf{r}\times\tau^B(\mathbf{u}^n - \mathbf{v}^n) d\Gamma \tag{21}$$

The last terms in Eqs. (20) and (21) are the penalty-like term that are added onto the surface force calculation. The total force acting on the object is the summation of the surface force and any external body forces (such as gravity and buoyancy).

4. Scalable immersogeometric analysis on octree meshes

While the IMGA is a promising approach to model particle migration in channels, it is challenging to optimize the computational efficiency on an unstructured background mesh. As a result, in this paper, we propose an octree-based IMGA that extends the IMGA on an octree-based background mesh. We employ the optimized parallel octree-based meshing library, Dendro. In this section, we discuss adaptations of IMGA as well as some computational aspects required and developed on Dendro that enable integration of IMGA on octree-based adaptive meshes.

4.1. Octree mesh implementation

While the elemental matrix computations are done using a separate external module (described in Section 4.2), Dendro provides the adaptive mesh refinement and all parallel data-structures. For this work, Dendro is extended to support meshes of (long) rectangular channels in order to account for non-cubic geometry domains. Octants outside the channels will be removed from the

[2] Due to the explicit time stepping used to track object motion, the $\Delta t$is usually set to a small value.

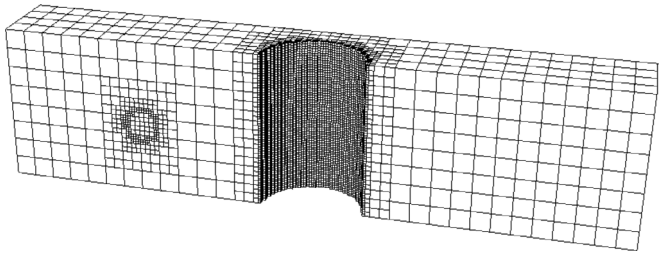

Fig. 3. An illustrative example of a rigid particle traversing a microchannel decorated with obstacles. Figure shows a slice cut through the geometry.

octree structure. Note that for channels with pillar obstacles, we can either immerse the pillars or create boundary-fitted structures since pillars will not move. The latter approach requires removal of octants inside pillars. An example of such an adaptively refined mesh (with a boundary-fitted pillar) is shown in Fig. 3.

The process in DENDRO used to build and maintain an adaptively refined octree mesh in parallel includes refinement, 2:1 balancing, partition, and meshing, and the algorithms of DENDRO are detailed in [55]. We also refer readers to [44,56–59] for details on implementation of DENDRO.

4.2. Elemental computation

Node coordinates and elemental connectivity are implicit in the octree's structure, so DENDRO recalculates these values on the fly as the octree is traversed. To avoid memory overhead, we consider a single hexahedral element. As we iterate through the octree mesh for assembly we re-position the nodes in this hexahedral element to match the octree element from DENDRO. Since the octree mesh has only one possible element shape, we pre-calculate and cache the isoparametric to physical mapping at each integration point. During initialization, we create an 'index' element at each refinement level in the octree and evaluate the basis functions at the integration points. When the assembly code needs to access these values, we pull them from the corresponding refinement level in the cache instead of recalculating them at each element.

4.3. Refinement according to in-out test and subdomains

To adapt the mesh refinement to the in-out test in IMGA, a coarse mesh is first constructed based on the geometry. Proceeding in a top-down fashion, each cell in the mesh is refined if a surface (pillar/particle) passes through it, which is determined using an in-out test. If all eight corners of an octant are outside of the immersed geometry, then we retain this element, but do not refine further. If all eight points are inside the immersed geometry, then this element is performed with the same manner as outside element for a immersed strategy, while it is removed from the octree for a boundary-fitted strategy. If some of the corners of the octant are inside and others outside, then this octant is refined. This process is repeated until the desired level of refinement is achieved. Similarily, the octants outside the rectangular channels are also removed by a channel boundary in-out test (as boundary-fitted strategy) during the refinement process. Channel boundaries may also be refined using the same way for a better approximation of the channel dimensions (and boundary layers if needed). Since in our case, the pillars, particles and channels are all regular geometries (i.e., cylinder, sphere and rectangular), the in-out test can be performed analytically. However for complex geometries, a ray-tracing algorithm may be employed in the in-out test.

In addition, subdomains, which leverage the original mesh datastructure, are created to handle meshes with rectangular geometry and holes for boundary-fitted pillars as octants outside the channel and inside the boundary-fitted pillars will be removed, and also no communications are needed for them. A different scattering mapping within the subdomains for current mesh is also uniquely defined afterwards. The finite element computations will only take place in subdomains (and we can discard the main octree structure for the original domain). Therefore, subdomains have an overall (much) smaller computation domain and store (significantly) less data than the original mesh (for example, in our case of a very long channel). Re-partitioning is required as creating subdomains will result in load imbalance. For our target application, it is important to identify both the external (channel) boundary as well as the internal boundary (boundary-fitted pillar surface). The subdomain stores two bits to keep track of whether a node is non-boundary, external, or internal boundary.

4.4. Sampling the immersed boundary and adding corrections

In order to reduce memory overhead and better parallelize the surface assembly in IMGA, we distribute surface quadrature points over processors. The object boundary mesh is generated as a triangulated mesh. Surface quadrature point coordinates, along with other necessary parameters, such as the unit normal vector and boundary values of velocity at each quadrature point, are then calculated in each triangle element using standard Gaussian quadrature. The surface quadrature points are then sorted and distributed over processes. This is done by associating each surface quadrature point with an octree element (real or virtual) with the maximum refinement that contains the quadrature point. Note that this octree element is not necessarily an existing octant in the octree mesh. This associated octree element represented by its bottom-left-back (minimum) node can be then aligned on the space-filling-curve, and the processor it belongs to can be easily found by the partitioning of the space-filling-curve. To find the actual background octree element that contains the quadrature point, we loop over all the octree elements in the process to check if the octree element is an ancestor of the associated octree element, or if they are exactly the same octree element. Since the octree elements and the surface quadrature points are both sorted based on the space-filling-curve, we can loop over them – in parallel – with an efficiency of $\mathcal{O}(m+n)$instead of a nested loop with an efficiency of $\mathcal{O}(m \times n)$,(unstructured meshes usually have to perform a nested loop), where $m$is the local number of elements in the background mesh and $n$is the local number of surface quadrature points. Boundary conditions imposed on the surface quadrature points can be then evaluated and distributed to the nodes of the background octree element. The distribution of surface quadrature points over processes and finding their background elements are challenging in unstructured meshes, as the process boundaries are usually complex in most graph-based partitioning approaches. When the object is moving, this is even more cumbersome since it has to be performed at each time step.

Another computational efficiency issue caused by the IMGA is that the immersed geometry is likely localized on a small subset of processes. These processes are the only ones that perform the surface assembly for weakly imposing no-slip boundary condition on the immersed boundary. A potential solution is to perform a weighted partition – increasing the weight of the intercepted elements by the additional relative cost of surface assembly with volume assembly. This weighted partition will ensure better load balancing. We defer this development to a subsequent paper.

4.5. Adaptive remeshing and intergrid transfers

An essential requirement for computational efficiency is to adapt the spatial mesh as the particle moves across the channel.

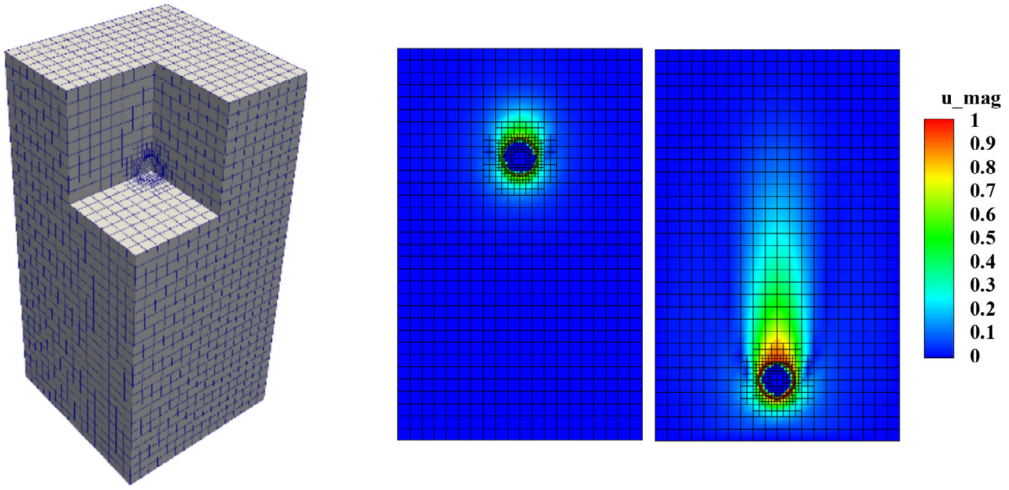

Fig. 4. A representative mesh illustrating the refinement around the particle, and contours of velocity magnitude at two representative time instances.

In the distributed memory setting, this also indicates a need to re-partition and re-balance the load. We adaptively remesh the domain at each time step based on the current position of the particle using results of the in-out test in IMGA followed by the subsequent 2:1 balance enforcement, partitioning and meshing process. Once the new mesh is generated, we transfer the data from the old mesh to the new mesh using interpolation. To keep things simple at this stage, we remesh each time from scratch followed by the repartition of new octree (reuse the same code of initial octree generation and partition because they are sufficiently optimized in DENDRO [59]), and then interpolate the local new mesh. The nodes of the local new mesh are distributed over processes based on the old mesh partitioning similarly as described in Section 4.4 to perform interpolation. Again, the intergrid transfer is challenging in unstructured meshes as repartitioning usually offers no guarantee of good overlap between the old and new partitions in most graph-based approaches, and the distribution of local new mesh will be difficult across complex process boundaries. A construction of global old mesh may be required for intergrid transfer in unstructured meshes.

5. Experiments and results

5.1. Implementation specification

The DENDRO framework is implemented in C++ using MPI for distributed memory parallelism and OpenMP for shared memory parallelism. This is integrated with a C++ module (Section 4.2) for evaluating basis functions and weak form of governing equations to support elemental computation. Our code is tightly integrated with PETSc v3.7's distributed matrix and vector data-structures and utilizes its SNES and KSP solvers. These tests were compiled and run on Oak Ridge's Titan supercomputer (before its decommissioning in 2019). PETSc, DENDRO, and the main program were compiled with the GNU 4.9.3 compiler with -O2 optimization flags. Timing information was reported using PETSc's logging framework.

5.2. Validation

We first validate the framework by comparing the particle trajectory and velocity against a benchmark experimental data of a sphere dropping in a quiescent fluid [60]. We consider a container with dimension of $0.1\,m \times 0.16\,m \times 0.1\,m$. We simulate a sphere with a diameter of $D = 0.015\,m$,released at a height of $0.12\,m$in the middle. The fluid has a density of $\rho_f = 960\,kg/m^3$,and a dynamic viscosity of $\mu = 0.058\,kg/(m \cdot s)$. The density of the sphere is $\rho_s = 1120\,kg/m^3$. Reynolds number, defined as $\frac{\rho_f u_0 D}{\mu}$,is $Re = 31.9$with a reference velocity $u_0 = 0.128\,m/s$. Time step size $\Delta t$is set to 1.2×10^{-3}s. Initial conditions are set as zero velocity in the whole fluid domain. No-slip boundary condition is imposed on lateral and bottom walls, and traction-free boundary condition is imposed on the top wall. We adaptively refine the mesh around the interface of the sphere and fluid. We refine three levels deeper than the rest of the background mesh. Specifically, we refine to a minimum/maximum level, $r = 5/8$(successively bisect and divide the octree root five and eight times, respectively). We remesh after each time step as the sphere drops. Note that such frequent adaptive remeshing is one of the challenges of our target application[3]. We set the surface triangular mesh size of the sphere in sync with the background interface element size, keeping a ratio of 1:2 (surface to background) to ensure adequate surface integration. The mesh example and visualizations of velocity magnitude contour, and the validation of non-dimensional height and sedimental velocity as the sphere settles downwards are presented in Figs. 4 and 5, respectively. As can be seen, both the sedimental velocity and the trajectory of the sphere match with the experiment results well.

5.3. Parallel scalability

We next show scaling performance of the framework. We collect timing for the case of a dropping sphere. We run each case for 5 time steps. The same setup and (re)meshing strategy as in last section is adopted. We run this experiment on four minimum/maximum refinement levels: $r = 5/8, 6/9, 7/10,$and $8/11$. Each refinement level has roughly seven to eight times more degrees of freedom to solve for than the previous level, with $r = 5/8$having 203,000 and $r = 8/11$reaching 70.2 million degrees of freedom.

We note that given specific minimum/maximum refinement level and the same initial and boundary conditions, the overall problem size (total degrees of freedom) in spite of remeshing is independent of the number of processes being used for the simulation. To this effect, we believe presenting performance for different minimum/maximum refinement levels with different numbers of processes, in the style of a strong scaling is appropriate. Indeed, performing weak scaling for such real-world applications

[3] While remeshing after every time step is not necessary, we perform this to illustrate scaling behavior of each part of the framework

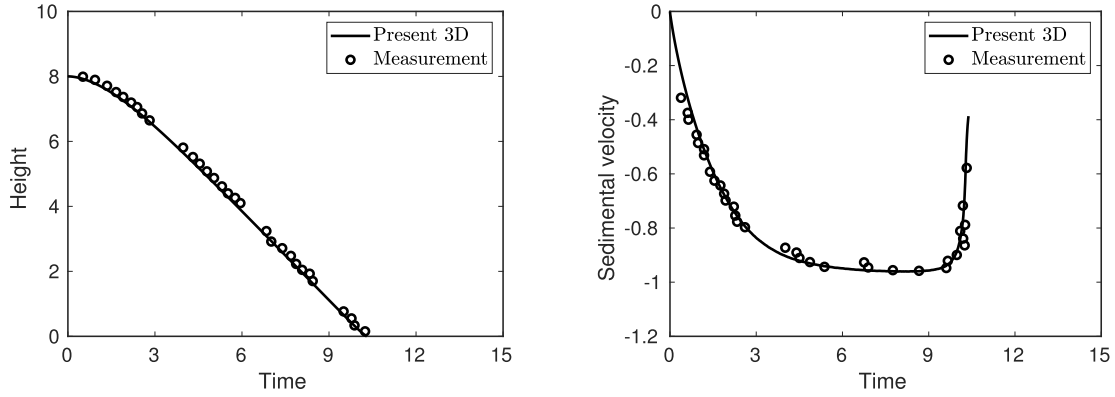

(a) Height vs. time. (b) Sedimental velocity vs. time.

Fig. 5. Comparisons of the non-dimensional height and sedimental velocity of the particle as it settles downwards with an experimental benchmark of a particle setting in a viscous fluid [60]. Note as the particle nears the bottom surface, its velocity rapidly zeros out.

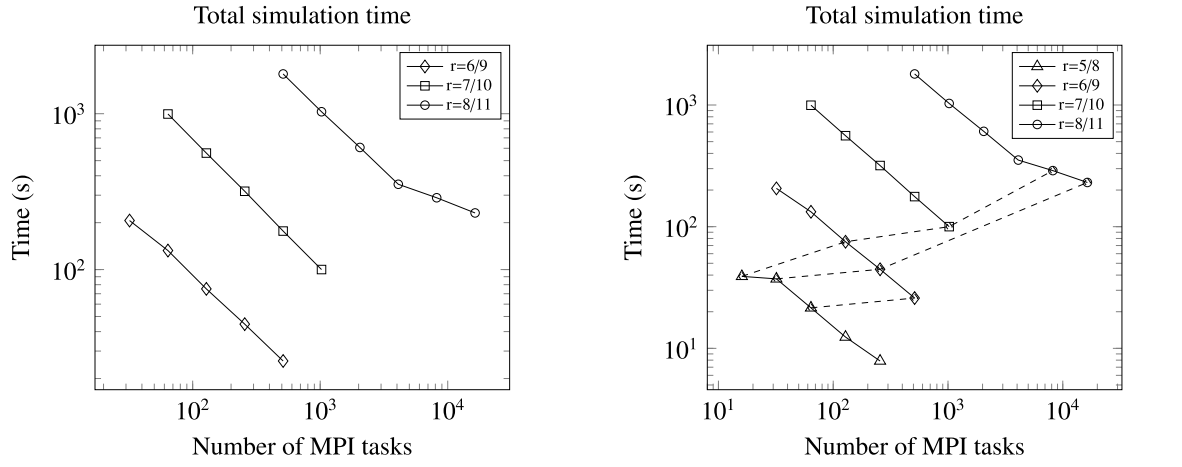

(a) Strong scaling for refinement level $r = 6/9, 7/10,$ and $8/11$

(b) Weak (dashed lines) scaling approximated from strong (solid lines) scaling results

Fig. 6. Strong and approximated weak scaling for a non-dimensional sphere of unit size dropping in a channel of size $8 \times 8 \times 8$running for 5 time steps with number of processes up to 16,384 on Titan.

is more difficult than strong scaling, since it is much harder to ensure that N/p.i.e., the grain size stays relatively constant with such frequent adaptive remeshing and consequently changes in problem size, where $N$is the total degrees of freedom and $p$is the number of processes. Therefore, given the somewhat fixed increase in problem size with increasing minimum/maximum refinement level and corresponding increase of number of processes, we can combine multiple strong scaling results to derive approximate weak scaling results for the overall simulation time. The approximated weak scaling results are presented in Fig. 6(b). Note that minor fluctuations in the approximation of the weak scalability are expected due to the inconsistent grain size.

5.3.1. Strong scalability

For our target application, the key goal is to be able to perform the simulations quickly. Given this, and the relatively moderate size of our problems, the focus is on strong scalability. We first present strong scalability results for the overall simulation time including the cost of everything in Fig. 6(a) for three problem sizes. Overall our code scales well, with continued reductions in simulation time. A breakdown of the total simulation time into various significant components for the refinement level of $r = 8/11$is also presented in Fig. 7. We can see that the amount of solve time and matrix assembly time, which are comparable, dominate the to-

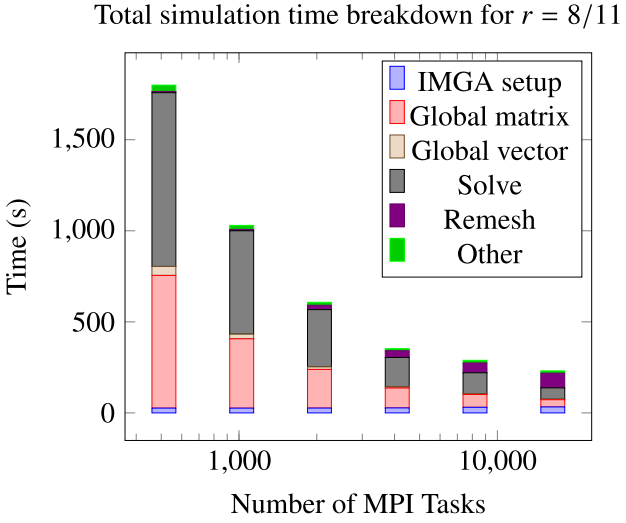

Fig. 7. Total simulation time broken down by category for refinement level $r = 8/11$. IMGA setup refers to the setup required to perform immersed boundary method. Global matrix and Global vector refer to the time taken to build the global Jacobian matrix and residual vector. Solve refers to the time taken to actually solve the system (i.e. PETScBCGS solver + ASM preconditioner). Remesh refers to the time taken to create the mesh of next time step and interpolate data onto it.

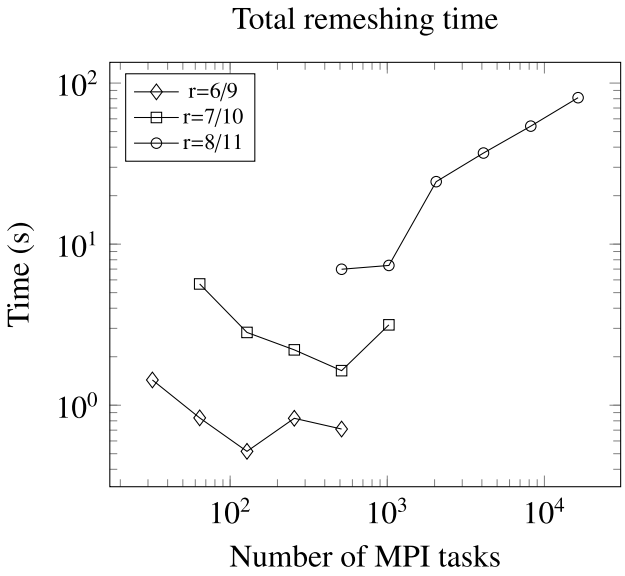

Fig. 8. Total time of adaptive remeshing for refinement level $r = 6/9, 7/10,$ and $8/11$.

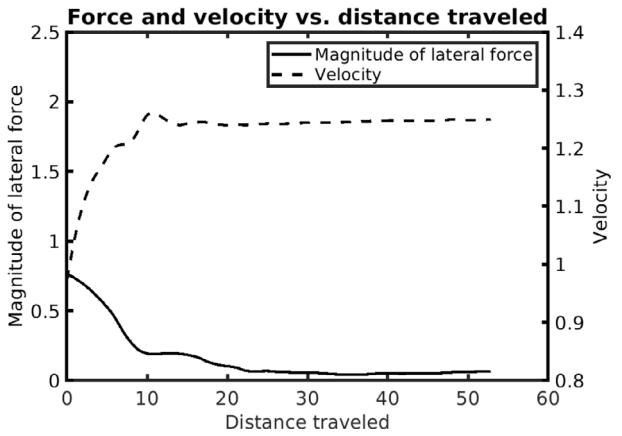

Fig. 10. Lateral force and velocity magnitude of the particle vs. distance it traveled downstream.

tal simulation time for most of the cases. The immersed boundary method corrections (IMGA setup) involving the distribution of immersed surface points on the octree mesh also scales reasonably well.

One significant trend with increasing number of processes is "total remeshing", as shown in Fig. 8 (also listed as "Remesh" in Fig. 7). This refers to the overall remeshing stage combining generating a new mesh, interpolating between two meshes and reinitializing the matrix, vector and solver. Effectively, this is the overhead paid for having good adaptivity. The scaling of remeshing is poor compared with other parts of the code, but the magnitude of time it takes is much smaller than solving the Navier–Stokes equations for most of the cases except using relatively large numbers of processes (last two data points) in the refinement level of $r = 8/11$. The remeshing time is comparable with the solve time (shown in Fig. 7) in these two cases when the communication becomes considerable. This is due to the interpolation between two meshes because the generation of a new adaptively refined mesh is sufficiently optimized. Note at the current stage, we perform a remeshing and repartitioning from scratch (due to the sufficiently optimized meshing code) first, followed by a subsequent interpolation. However, the interpolation may not be optimal since a large amount of communication may be required by distributing new local mesh. We are exploring alternatives (to be reported in a subsequent paper). Specifically, we could remesh (refine or coarsen octants) in each process while keeping the local geometry domain unchanged in each process. This means the old and new local mesh are overlapping and consequently there is no need to distribute new local mesh over processes during interpolation. We could then perform interpolation (no communication needed) first in each process, and then repartition the new octree and corresponding newly interpolated solution vector for load balancing.

5.4. Results for particle tracking in microchannels

We finally present two results for the application of this framework. The first is particle tracking in a channel with a square cross-section, and the other is our canonical problem of particle tracking in a channel with pillar obstacles. The schematics of both cases are shown in Fig. 9, and we present both cases in non-dimensional units.

5.4.1. Square channel

Case setup: We consider a long channel with dimensions $96 \times 4 \times 4$. We simulate a spherical particle released at$(10, 2, 1)$ with diameter of $D = 1$. The origin is located at the bottom-left-front corner (oriented as shown in Fig. 9) of the channel. We set the particle Reynolds number, $Re_D = 5$. We assume the particle is of the same density as the fluid, so that there is no buoyancy effect. Time step size $\Delta t$is set to 0.05. The inlet has unit velocity normal to the inlet. No-slip boundary condition is imposed on surrounding walls, and zero pressure is imposed on the outlet. The initial condition for the fluid velocity as well as the particle velocity are both set to be the same as the inlet velocity. We note that such long simulations – tracking the temporal evolution of the particle as it traverses nearly 50D downstream – is fairly atypical in the microfluidics community.

Results: We are interested in the magnitude of lateral force acting on the particle and the magnitude of the particle velocity as the particle reaches its equilibrium cross-sectional position as shown in Fig. 10. After the particle has traveled 30Ddownstream,

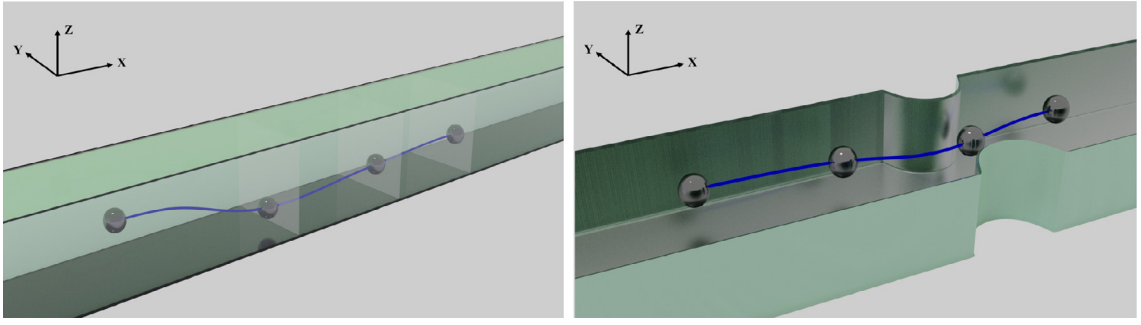

Fig. 9. Schematics of particle tracking in different configurations.

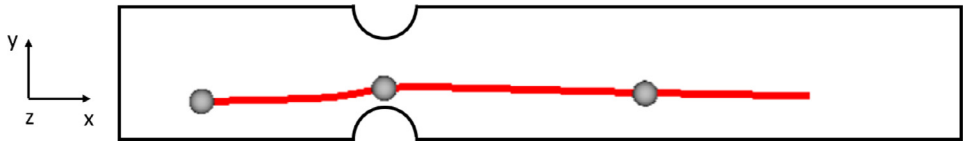

Fig. 11. Top-down view ($x-y$plane) of pathline of particle movement in a channel with pillars. Boundary-fitted pillars are plotted for visualization purpose while they are immersed in the actual simulation.

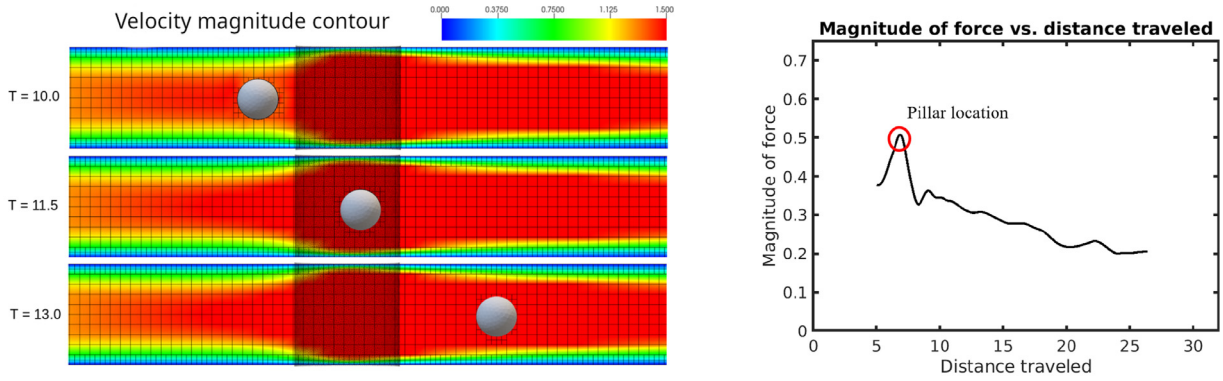

(a) Side view ($x-z$ plane) of flow velocity magnitude in channel with pillars at different time steps. The mesh shown is sub-sampled from the actual mesh for ease of visualization purpose.

(b) Magnitude of force on the particle vs. distance it traveled downstream.

Fig. 12. Velocity magnitude of particle in channel with pillars at different time steps, and force magnitude on the particle during its motion.

it is clear that the velocity have converged to a steady value (representing pure streamwise motion, and very small lateral motion), which suggests that the particle has reached its equilibrium position. Additionally, the net lateral force becomes negligibly small. The final cross sectional location in $y-z$plane is (2.0, 1.04) which matches the experimentally determined equilibrium position [61].

5.4.2. Channel with pillars

Case setup: In our canonical problem, we consider a sphere of diameter of $D = 1$in a channel of dimensions of $32 \times 5 \times 2.5$. Two half pillars of radius 1.25 and height 2.5 are placed in the channel, forming a converging-diverging type of cross section. The particle is released from $(3, 1.4, 1.25)$ with the same placement of the origin as the previous example. The particle Reynolds number is $Re_D = 50$using the channel flow rate, which ensures that inertial effects are prominent [3]. Time step size $\Delta t$is set to 0.015. The boundary conditions are the same as the previous case except for the two half pillars. Note, we also immerse the two half pillars, and therefore the no-slip boundary condition on the pillars are weakly enforced. The initial condition of fluid velocity is the same as inlet velocity. The sphere is held stationary until $t = 5$to wait for the channel flow to fully develop, so that a physically meaningful force is imposed on the released particle.

Results: Fig. 11 plots the particle path as it navigates the channel with the pillar obstacles. This path curves in as the particle passes the pillars, and due to the inertial regime that the flow is in, the cross-sectional position of the particle close to outlet is offset from the initial cross-sectional location. This is in line with expected behavior from experiments in Stoecklein and Di Carlo [3], which suggest that inertial microfluidics with pillars can produce irreversible cross-sectional displacements.

Fig. 12 (a) illustrates the 'squeezing' effect due to the presence of pillars and plots the flow velocity magnitude contour along the $x-z$plane (i.e., side view) at 3 different time steps (before, during and after the particle interacts with the pillars). Note that there is no direct interaction between the particle and the pillars, but instead all interactions are mediated by the fluid. Fig. 12(b) quantifies this observation by plotting the force acting on the particle. Note the large jump in force as the particle traverses the channel (close to the pillar) is due to the squeezing effect. Furthermore, the simulation was performed on a mesh with around 105,000 hexahedra elements, the average number of degrees of freedom for this problem is around 350K, and the number of time steps needed to track the particle across the channel dimension is 790 steps. The total simulation time for this canonical problem is around 10 hours using 12 KNL nodes on TACC Stampede2. This is very promising as it allows us to proceed with computing cross-sectional displacement maps under different pillar configurations, which essentially translates to executing this type of simulation for a large set of different release locations across the inlet cross-section.

6. Conclusions and future directions

We developed a scalable, adaptively refined octree-based immersogeometric analysis framework, and validated this framework using a benchmark case of sphere dropping in fluids. Our framework demonstrates excellent strong (and weak) scalability for the overall simulation time, even with frequent remeshing in the benchmark case. Our framework can keep the overhead of adaptive remeshing and IMGA corrections relatively low. We anticipate additional code optimization will make the approach even more scalable. We further deployed the framework to track particle in microchannels with different (complex) geometries. This framework allows us to efficiently construct the deformation maps for particles under a broad range of experimentally accessible parameters, which will result in a passive approach for particle localization. We identify several avenues of future work. Immediate computational goals include (1) transitioning to a matrix-free solver that can significantly reduce the solve-time, while ensuring sustained adaptivity for larger processor counts, (2) designing a local refinement/coarsening algorithm and subsequent repartitioning to optimize intergrid transfer, (3) incorporating a more rigorous dynamic

load balancing that accounts for the additional work involved in the surface computations, and (4) accounting for multiple moving objects. From the flow physics perspective, we plan to deploy this framework to characterize the inertial displacement maps for a range of particle sizes and pillar placements.

Declaration of Competing Interest

We confirm that this work is original and has not been published elsewhere, nor is it currently under consideration for publication elsewhere. We have no conflicts of interest to disclose.

Acknowledgment

The authors acknowledge XSEDE grant number TG-CTS110007 for computing time on TACC Stampede2, Oak Ridge's Titan supercomputing resource and Iowa State University computing resources. Ganapathysubramanian and Gao were funded in part by NSF grant 1935255.

References

[1] Stoecklein D, Di Carlo D. Nonlinear microfluidics. Anal Chem 2019;91(1):296–314.
[2] Gossett DR, Weaver WM, Mach AJ, Hur SC, Tse HTK, Lee W, et al. Label-free cell separation and sorting in microfluidic systems. Anal Bioanal Chem 2010;397(8):3249–67.
[3] Stoecklein D, Di Carlo D. Nonlinear microfluidics. Anal Chem 2018;91(1):296–314.
[4] Peskin CS. Flow patterns around heart valves: a numerical method. J Comput Phys 1972;10(2):252–71.
[5] Peskin CS. Flow patterns around heart valves: a digital computer method for solving the equations of motion. IEEE Trans Biomed Eng 1973(4):316–17.
[6] Glowinski R, Pan T-W, Hesla TI, Joseph DD. A distributed Lagrange multiplier/fictitious domain method for particulate flows. Int J Multiphase Flow 1999;25(5):755–94.
[7] Glowinski R, Pan TW, Hesla TI, Joseph DD, Périaux J. A fictitious domain approach to the direct numerical simulation of incompressible viscous flow past moving rigid bodies: application to particulate flow. J Comput Phys 2001;169(2):363–426.
[8] Glowinski R, Kuznetsov Y. Distributed Lagrange multipliers based on fictitious domain method for second order elliptic problems. Comput Methods Appl Mech Eng 2007;196:1498–506.
[9] Zhang L, Gerstenberger A, Wang X, Liu WK. Immersed finite element method. Comput Methods Appl Mech Eng 2004;193:2051–67.
[10] Liu WK, Kim DW, Tang S. Mathematical foundations of the immersed finite element method. Comput Mech 2007;39(3):211–22.
[11] Wang XS, Zhang LT, Liu WK. On computational issues of immersed finite element methods. J Comput Phys 2009;228(7):2535–51.
[12] Wang X, Zhang LT. Modified immersed finite element method for fully-coupled fluid–structure interactions. Comput Methods Appl Mech Eng 2013;267:150–69.
[13] Casquero H, Bona-Casas C, Gomez H. A NURBS-based immersed methodology for fluid–structure interaction. Comput Methods Appl Mech Eng 2015;284:943–70.
[14] Kamensky D, Hsu M-C, Schillinger D, Evans JA, Aggarwal A, Bazilevs Y, et al. An immersogeometric variational framework for fluid–structure interaction: application to bioprosthetic heart valves. Comput Methods Appl Mech Eng 2015;284:1005–53.
[15] Xu F, Schillinger D, Kamensky D, Varduhn V, Wang C, Hsu M-C. The tetrahedral finite cell method for fluids: immersogeometric analysis of turbulent flow around complex geometries. Comput Fluids 2016;141:135–54.
[16] Hsu M-C, Wang C, Xu F, Herrema AJ, Krishnamurthy A. Direct immersogeometric fluid flow analysis using B-rep CAD models. Comput Aided Geom Des 2016;43:143–58.
[17] Wang C, Xu F, Hsu M-C, Krishnamurthy A. Rapid B-rep model preprocessing for immersogeometric analysis using analytic surfaces. Comput Aided Geom Des 2017;52–53:190–204.
[18] Xu F, Bazilevs Y, Hsu M-C. Immersogeometric analysis of compressible flows with application to aerodynamic simulation of rotorcraft. Math Models Methods Appl Sci 2019;29:905–38.
[19] Bazilevs Y, Hughes TJR. Weak imposition of Dirichlet boundary conditions in fluid mechanics. Comput Fluids 2007;36:12–26.
[20] Xu S, Xu F, Kommajosula A, Hsu M-C, Ganapathysubramanian B. Immersogeometric analysis of moving objects in incompressible flows. Comput Fluids 2019;189:24–33.
[21] Bazilevs Y, Calo VM, Cottrel JA, Hughes TJR, Reali A, Scovazzi G. Variational multiscale residual-based turbulence modeling for large eddy simulation of incompressible flows. Comput Methods Appl Mech Eng 2007;197:173–201.
[22] Xu S, Gao B, Hsu M-C, Ganapathysubramanian B. A residual-based variational multiscale method with weak imposition of boundary conditions for buoyancy-driven flows. Comput Methods Appl Mech Eng 2019;352:345–68.
[23] Xu S, Liu N, Yan J. Residual-based variational multi-scale modeling for particle-laden gravity currents over flat and triangular wavy terrains. Comput Fluids 2019;188:114–24.
[24] Bazilevs Y, Takizawa K, Tezduyar TE, Hsu M-C, Kostov N, McIntyre S. Aerodynamic and FSI analysis of wind turbines with the ALE–VMS and ST–VMS methods. Arch Comput Methods Eng 2014;21:359–98.
[25] Yan J, Lin S, Bazilevs Y, Wagner GJ. Isogeometric analysis of multi-phase flows with surface tension and with application to dynamics of rising bubbles. Comput Fluids 2018.
[26] Zhu Q, Xu F, Xu S, Hsu M-C, Yan J. An immersogeometric formulation for free–surface flows with application to marine engineering problems. Comput Methods Appl Mech Eng 2020;361:112748.
[27] Yan J, Deng X, Xu F, Xu S, Zhu Q. Numerical simulations of two back–to-back horizontal axis tidal stream turbines in free-surface flows. J Appl Mech 2020;87(6).
[28] Takizawa K, Tezduyar TE, Kuraishi T. Multiscale space–time methods for thermo-fluid analysis of a ground vehicle and its tires. Math Models Methods Appl Sci 2015;25(12):2227–55.
[29] Sondak D, Shadid JN, Oberai AA, Pawlowski RP, Cyr EC, Smith TM. A new class of finite element variational multiscale turbulence models for incompressible magnetohydrodynamics. J Comput Phys 2015;295:596–616.
[30] Liu J, Oberai AA. The residual-based variational multiscale formulation for the large eddy simulation of compressible flows. Comput Methods Appl Mech Eng 2012;245:176–93.
[31] Xu S. Buoyancy-driven flow and fluid-structure interaction with moving boundaries. Iowa State University; 2018. Graduate Theses and Dissertations. https://lib.dr.iastate.edu/etd/17364.
[32] Lofquist A. A scalable software framework for solving PDEs on distributed octree meshes using finite element methods. Iowa State University; 2018. Graduate Theses and Dissertations. https://lib.dr.iastate.edu/etd/16842.
[33] Bielak J, Ghattas O, Kim EJ. Parallel octree-based finite element method for large-scale earthquake ground motion simulation. Comput Model Eng Sci 2005;10(2):99.
[34] Becker R, Braack M. Multigrid techniques for finite elements on locally refined meshes. Numer Linear Algebra Appl 2000;7(6):363–79.
[35] Legrain G, Allais R, Cartraud P. On the use of the extended finite element method with quadtree/octree meshes. Int J Numer Methods Eng 2011;86(6):717–43.
[36] Thieulot C, Fullsack P, Braun J. Adaptive octree-based finite element analysis of two-and three-dimensional indentation problems. J Geophys Res 2008;113(B12).
[37] Patra AK, Laszloffy A, Long J. Data structures and load balancing for parallel adaptive hp finite-element methods. Comput Math Appl 2003;46(1):105–123.
[38] Flaherty JE, Loy RM, Özturan C, Shephard MS, Szymanski BK, Teresco JD, et al. Parallel structures and dynamic load balancing for adaptive finite element computation. Appl Numer Math 1998;26(1–2):241–63.
[39] Fernando M, Neilsen D, Lim H, Hirschmann E, Sundar H. Massively parallel simulations of binary black hole intermediate-mass-ratio inspirals. SIAM J Sci Comput 2019;41(2):C97–C138.
[40] Fernando M, Neilsen D, Hirschmann EW, Sundar H. A scalable framework for adaptive computational general relativity on heterogeneous clusters. In: Proceedings of the ACM international conference on supercomputing. ACM; 2019. p. 1–12.
[41] Ishii M, Fernando M, Saurabh K, Khara B, Ganapathysubramanian B, Sundar H. Solving PDEs in space-time: 4D tree-based adaptivity, mesh-free and matrix-free approaches. In: SC'19: Proceedings of the international conference for high performance computing, networking, storage and analysis. ACM; 2019. 61:1–61:61.
[42] Sundar H, Davatzikos C, Biros G. Biomechanically-constrained 4d estimation of myocardial motion. In: Medical image computing and computer-assisted intervention – MICCAI 2009. Springer Berlin Heidelberg; 2009. p. 257–65.
[43] Sundar H, Shen D, Biros G, Xu C, Davatzikos C. Robust computation of mutual information using spatially adaptive meshes. In: Proceedings of the 10th international conference on medical image computing and computer-assisted intervention - Volume Part I. Springer-Verlag; 2007. p. 950–8.
[44] Sundar H, Sampath RS, Adavani SS, Davatzikos C, Biros G. Low-constant parallel algorithms for finite element simulations using linear octrees. In: SC'07: Proceedings of the international conference for high performance computing, networking, storage, and analysis. ACM/IEEE; 2007.
[45] Sundar H, Sampath RS, Biros G. Bottom-up construction and 2: 1 balance refinement of linear octrees in parallel. SIAM J Sci Comput 2008;30(5):2675–708.
[46] Tu T, O'Hallaron DR, Ghattas O. Scalable parallel octree meshing for terascale applications. In: Proceedings of the 2005 ACM/IEEE conference on Supercomputing. IEEE Computer Society; 2005. p. 4.
[47] Warren MS, Salmon JK. A parallel hashed oct-tree N-body algorithm. In: Supercomputing '93 proceedings; 1993. p. 12–21.
[48] Ying L, Biros G, Zorin D, Langston H. A new parallel kernel-independent fast multipole method. In: SC'03: Proceedings of the 2003 ACM/IEEE conference on supercomputing. IEEE; 2003. p. 14.
[49] Baraff D. An introduction to physically based modeling: rigid body simulation i–unconstrained rigid body dynamics. In: SIGGRAPH Course Notes; 1997. p. D31–68.

[50] Nitsche J. Über ein Variationsprinzip zur Lösung von Dirichlet-Problemen bei Verwendung von Teilräumen, die keinen Randbedingungen unterworfen sind. Abhandlungen aus dem Mathematischen Seminar der Universität Hamburg 1971;36:9–15.
[51] Bazilevs Y, Michler C, Calo VM, Hughes TJR. Weak Dirichlet boundary conditions for wall-bounded turbulent flows. Comput Methods Appl Mech Eng 2007;196:4853–62.
[52] Bazilevs Y, Michler C, Calo VM, Hughes TJR. Isogeometric variational multiscale modeling of wall-bounded turbulent flows with weakly enforced boundary conditions on unstretched meshes. Comput Methods Appl Mech Eng 2010;199:780–90.
[53] Wu MCH, Kamensky D, Wang C, Herrema AJ, Xu F, Pigazzini MS, et al. Optimizing fluid–structure interaction systems with immersogeometric analysis and surrogate modeling: application to a hydraulic arresting gear. Comput Methods Appl Mech Eng 2017;316:668–93.
[54] Balay S, Abhyankar S, Adams MF, Brown J, Brune P, Buschelman K, et al. PETSc users manual Tech. Rep. ANL-95/11 - Revision 3.10. Argonne National Laboratory; 2018. http://www.mcs.anl.gov/petsc.
[55] Sundar H, Sampath R, Biros G. Bottom-up construction and 2:1 balance refinement of linear octrees in parallel. SIAM J Sci Comput 2008;30(5):2675–708.
[56] Bern M, Eppstein D, Teng S-H. Parallel construction of quadtrees and quality triangulations. Int J Comput Geom.Appl. 1999;9(06):517–32.
[57] Burstedde C, Wilcox LC, Ghattas O. p4est: scalable algorithms for parallel adaptive mesh refinement on forests of octrees. SIAM J Sci Comput 2011;33(3):1103–33.
[58] Sundar H, Biros G, Burstedde C, Rudi J, Ghattas O, Stadler G. Parallel geometric-algebraic multigrid on unstructured forests of octrees. In: SC'12: Proceedings of the international conference on high performance computing, networking, storage and analysis. Los Alamitos, CA, USA: IEEE Computer Society Press; 2012 http://dl.acm.org/citation.cfm?id=2388996.2389055. ISBN 978-1-4673-0804-5. 43:1–43:11.
[59] Fernando M, Duplyakin D, Sundar H. Machine and application aware partitioning for adaptive mesh refinement applications. In: Proceedings of the 26th international symposium on high-performance parallel and distributed computing. ACM; 2017. p. 231–42.
[60] Ten Cate A, Nieuwstad CH, Derksen JJ, Van den Akker HEA. Particle imaging velocimetry experiments and lattice-Boltzmann simulations on a single sphere settling under gravity. Phys Fluids 2002;14(11):4012–25.
[61] Di Carlo D, Edd JF, Humphry KJ, Stone HA, Toner M. Particle segregation and dynamics in confined flows. Phys Rev Lett 2009;102(9):094503.