ELSEVIER

Curvilinear immersed boundary method for simulating fluid structure interaction with complex 3D rigid bodies

Iman Borazjani [a], Liang Ge [b], Fotis Sotiropoulos [a],*

[a] *St. Anthony Falls Laboratory, University of Minnesota, 2 Third Avenue SE, Minneapolis, MN 55414, United States*
[b] *Department of Surgery, University of California, San Francisco, CA 94143, United States*

Abstract

The sharp-interface CURVIB approach of Ge and Sotiropoulos [L. Ge, F. Sotiropoulos, A numerical method for solving the 3D unsteady incompressible Navier–Stokes equations in curvilinear domains with complex immersed boundaries, Journal of Computational Physics 225 (2007) 1782–1809] is extended to simulate fluid structure interaction (FSI) problems involving complex 3D rigid bodies undergoing large structural displacements. The FSI solver adopts the partitioned FSI solution approach and both loose and strong coupling strategies are implemented. The interfaces between immersed bodies and the fluid are discretized with a Lagrangian grid and tracked with an explicit front-tracking approach. An efficient ray-tracing algorithm is developed to quickly identify the relationship between the background grid and the moving bodies. Numerical experiments are carried out for two FSI problems: vortex induced vibration of elastically mounted cylinders and flow through a bileaflet mechanical heart valve at physiologic conditions. For both cases the computed results are in excellent agreement with benchmark simulations and experimental measurements. The numerical experiments suggest that both the properties of the structure (mass, geometry) and the local flow conditions can play an important role in determining the stability of the FSI algorithm. Under certain conditions the FSI algorithm is unconditionally unstable even when strong coupling FSI is employed. For such cases, however, combining the strong coupling iteration with under-relaxation in conjunction with the Aitken's acceleration technique is shown to effectively resolve the stability problems. A theoretical analysis is presented to explain the findings of the numerical experiments. It is shown that the ratio of the added mass to the mass of the structure as well as the sign of the local time rate of change of the force or moment imparted on the structure by the fluid determine the stability and convergence of the FSI algorithm. The stabilizing role of under-relaxation is also clarified and the upper bound of the under-relaxation coefficient, required for stability, is derived.
© 2008 Elsevier Inc. All rights reserved.

Keywords: Incompressible flow; Fluid structure interaction; Immersed boundaries; Heart valves

---

* Corresponding author.
 E-mail address: fotis@umn.edu (F. Sotiropoulos).

1. Introduction

Fluid structure interaction (FSI) problems are encountered in a wide range of engineering areas, including, among others, aeronautical engineering, coastal engineering, and biomedical engineering. FSI problems of practical interest often involve multiple 3D immersed bodies with complex geometry undergoing very large structural displacements and inducing very complex flow phenomena. Consider for example a bileaflet mechanical heart valve (BMHV), which is currently the most widely implanted prosthesis for replacing diseased natural heart valves. A BMHV consists of two leaflets each of which is allowed to rotate freely (see Fig. 1 for an illustration of the bileaflet geometry) such that when surgically implanted in a patient the leaflets can open and close under the action of the pulsatile blood flow induced by the contracting and expanding heart chambers, respectively. During the opening and closing phases the valve leaflets undergo large displacement, sweeping an angle of about 55° (Fig. 1) and experience very large acceleration rates [1]. Recent experiments and simulations with prescribed leaflet kinematics [1,2] have shown that the flow at physiologic conditions (the Reynolds number at peak systole can be as high as $Re = 6000$, based on the bulk inlet flow and the aorta diameter) is very complex especially past peak systole when the flow through the valve decelerates and breaks down rapidly into a chaotic, turbulent-like state. The robust, accurate and efficient numerical simulation of such a complex FSI problem involving large structural displacements and very complex flow phenomena is a good representative of the types of real-life engineering problems we alluded to at the beginning of this paragraph and serves as the main motivation of our work.

A FSI algorithm consists of three major numerical components: (1) the algorithm for solving the governing equations in a domain with moving immersed boundaries; (2) the iterative strategy for coupling together the fluid and structural domains; and (3) the algorithm for efficiently solving the grid connectivity problem, i.e., for identifying the instantaneous location of the moving bodies within the flow domain. In the remainder of this section we review the related literature for each one of these three components, identify the objectives and contributions of our work, and present the outline of the paper.

Numerical techniques for simulating flows with moving boundaries can be classified in two broad categories: (1) moving grid methods, mainly the arbitrary Lagrangian–Eulerian (ALE) approach [3]; and (2) fixed grid methods, such as the immersed boundary method developed by [4]. ALE methods employ a grid that is adapted to and moves and deforms with the moving boundary. Such methods have been applied, among others, to study the transient aeroelastic response of airfoils [5], the FSI problem of a shock absorber [6], the blood flow through compliant aortas [7], etc. A significant limitation of the ALE approach, however, stems from the fact that the mesh conforms to the moving boundary and as such it needs to be constantly displaced and deformed following the motion of the boundary. The mesh moving step could be quite challenging and expensive for complicated 3D problems. This situation is further exacerbated in problems involving large structural displacements for which frequent remeshing might be the only feasible approach for ensuring a well conditioned mesh at each time step of the simulation. Because of this inherent limitation, the ALE approach is only applicable to FSI problems involving relatively small structural displacements. In fixed grid approaches, on the other hand, the entire computational domain (including both the fluid and structure domains) is discretized with a single, fixed, non-boundary conforming grid system (most commonly a Cartesian mesh is used
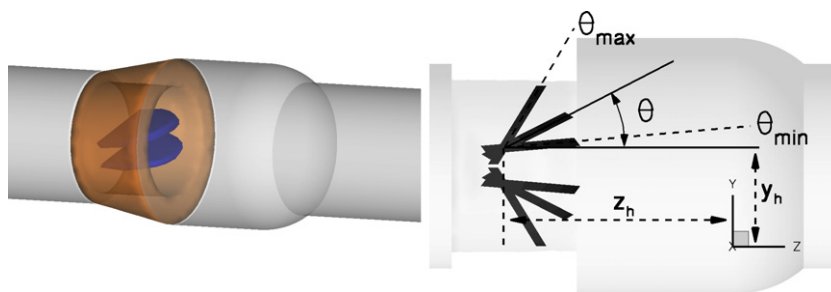

Fig. 1. The 3D BMHV geometry including the housing and the leaflets in a straight aorta (left) and the definition of leaflet angle (right). The leaflets hinge axis is placed parallel to x-axis of the global frame of reference at position (y_h, z_h) in the YZ plane.

as the fixed background mesh). The effect of a moving immersed boundary is accounted for by adding forcing terms to the governing equations of fluid motion so that the presence of a no-slip boundary at the location of the interface can be felt by the surrounding flow. Because of the fixed grid arrangement, such methods are inherently applicable to FSI problems involving arbitrarily large structural displacements.

The classical representative of fixed grid methods is Peskin's immersed boundary method, which has been successfully applied, among others, to FSI problems arising in heart valve fluid mechanics [4], biofilming processes [8], the interaction between flexible fibers and flow [9], the interaction between a flapping filament and ambient flow [10], etc. Under the original immersed boundary framework, however, the effects of the immersed boundaries on the fluid flow is accounted for through a discrete delta function, which results in the artificial spreading or diffusion of the boundary interface across a number of grid points. Because of this property, immersed boundary methods are known as diffused interface methods and typically require increased grid resolution in the vicinity of the boundary for accurate results especially for problems involving high Reynolds numbers. To remedy this situation, a class of sharp-interface immersed boundary methods has recently been developed and are attracting increasing attention in the literature – see for example [11–17] among many others and some recent publications from our group [2,18,19]. In these methods the immersed boundary is treated as a sharp interface and its effect on the fluid is accounted for either by locally modifying the shapes of the grid cells to produce a boundary-fitted mesh (cut-cell methods [20]), or by incorporating the jump conditions caused by the discrete delta function into the finite difference scheme to avoid the approximation of the discrete delta function by a smooth function (immersed interface methods [16,17,21]), or by reconstructing boundary conditions in the vicinity of the immersed boundary via an appropriate interpolation scheme (hybrid Cartesian/immersed boundary (HCIB) methods [19]). The reader is referred to [19] and the recent review paper by Mittal and Iaccarino [22] for a detailed discussion of various sharp-interface methodologies. In a previous publication from our group we have shown the great promise of HCIB methods by reporting results for various flows with complex moving boundaries, including flows induced by flapping wings, swimming fishes and planktonic copepods [19]. In a more recent publication [2] we proposed a novel modeling paradigm which integrates the HCIB approach with curvilinear, body-fitted meshes. This method, which we dubbed the Curvilinear immersed boundary (CURVIB) method, is very well suited for problems for which the moving boundary is immersed in a domain that can be discretized efficiently with a curvilinear body-fitted mesh. The CURVIB method was successfully applied to perform the first direct numerical simulation of physiologic pulsatile flow in a BMHV mounted in a straight aorta (see [2] for the details of the method and [1] for validation and discussion of BMHV flow physics). It should be noted, however, that both in [1,2] the simulations were carried out by prescribing the kinematics of the valve leaflets from experimental measurements and as such the FSI aspects of the BMHV problem were not considered. It is the main objective of this work to extend the CURVIB approach of Ge and Sotiropoulos [2] to develop a coupled FSI formulation that is applicable to problems involving multiple, moving, rigid bodies of complex geometry undergoing arbitrarily large structural displacements. To the best of our knowledge, the only FSI method using a sharp-interface Cartesian framework that is somewhat related to our present effort was recently reported in [23]. In this work, however, the governing equations are formulated in the reference frame moving with the immersed body and as such the method is difficult to apply to problems involving multiple moving bodies.

In a FSI problem the fluid and structure motions are governed by different albeit tightly coupled sets of equations. A monolithic approach that solves the two sets of governing equations simultaneously is usually very expensive and currently the most widely adopted FSI solution strategy is the so-called partitioned approach, in which the fluid and structure domains are treated as two separate domains. That is, the governing equations of each domain are considered to be independent of one another and are coupled together through matching boundary conditions at the interfaces. Currently there are two coupling strategies, namely the so-called *loose coupling* and *strong coupling* – in this paper we shall refer to these coupling strategies as LC-FSI and SC-FSI, respectively. Examples of LC-FSI can be found in [24–26] while SC-FSI implementations can be found in [27,28], among others. As illustrated in Section 2.2 below, a major difference between LC-FSI and SC-FSI is that they respectively integrate the governing equations of the structure explicitly and implicitly in time. The implicit integration in the SC-FSI approach is accomplished by carrying out a number of sub-iteration steps until the coupled fluid/structure solution converges to a stable solution at every physical

time step. This major difference between the two coupling strategies leads to inherently different numerical properties. Namely, LC-FSI is considerably less stable and robust than its strong coupling counterpart but requires much lower computational cost per time step as no sub-iterations are required. In the current paper, we implement both the loose coupling and strong coupling FSI strategies in order to elucidate and quantify the differences between these two methods through a series of numerical experiments reported in Section 3.

An important aspect of sharp-interface immersed boundary methods, and all fixed grid approaches, that becomes especially critical in FSI implementations where the motion of the immersed body is not known a priori, is the issue of interface tracking. Since the immersed objects are detached from and move within the background grid one needs an efficient scheme for tracking the motion of the interfaces and establishing the relationship between the immersed bodies and the background grid at every time step of the FSI simulation. Existing interface tracking schemes can be loosely classified in two categories: (a) front tracking approaches; and (b) front capturing approaches. In a front tracking approach, the interface surfaces are discretized with a set of Lagrangian markers whose trajectories are explicitly tracked in a Lagrangian manner. Applications of front tracking methods can be found in [20,29,19], among others. In the front capturing approach, on the other hand, the interface tracking is usually accomplished by introducing an auxiliary scalar variable defined on the fixed grid system and reconstructing the location and shape of the interfaces using the contours or level-sets of this scalar variable. Variations of this approach include the level-set method [30], the volume of fluid (VOF) method [31], etc. The front tracking approach is inherently more accurate, due to the explicitly Lagrangian tracking of the interfaces, than the front capturing approach, which typically introduces some numerical diffusion at the interface. The implementation of the front tracking approach, however, is much more complicated than front capturing approaches in complex 3D problems. In this work we adopt the front tracking method because its inherent support of precise interface description, which is a strongly desirable feature in FSI simulations involving multiple complex bodies. To facilitate the efficient implementation of the front tracking approach, we develop an efficient numerical procedure that can quickly identify the relationship between the background Eulerian grid and the Lagrangian grids describing the interfaces – or equivalently identify whether a background grid node is in the interior or exterior of the immersed objects. The algorithm is based on the idea of ray-tracing algorithm [32] and is an extension to 3D of the efficient 2D point-in-polygon algorithm, which is used to identify whether a point in a 2D plane is located inside or outside of a polygon that is discretized by a set of 2D line segments. We will show that the new 3D grid relationship identification algorithm is very efficient and is capable to handle arbitrarily complex immersed bodies.

This paper is organized as follows: In Section 2 we discuss in detail all aspects of our FSI numerical method. This includes the governing equations, the loose coupling and strong coupling FSI solver, and the efficient front tracking method with fast grid relationship identification procedure. In Section 3 we report results for two FSI problems in order to validate and illustrate the capabilities of the solver: (a) vortex induced vibration (VIV) of 2D elastically mounted circular cylinders; and (b) pulsatile flow through a clinical quality BMHV at physiologic conditions. In the same section we also explore via numerical experiments the stability and robustness of the LC-FSI and SC-FSI algorithms. In Section 3.2.3 we present a simple theoretical analysis to explain the findings of these numerical experiments and identify key structural, flow-related, and numerical parameters that govern the stability of FSI algorithms. Finally we summarize the major findings of this work and identify areas for further research.

2. Numerical method

In this section we present all aspects of our numerical approach. We begin by presenting the governing equations of the partitioned FSI approach. Subsequently we discuss the numerical algorithm for coupling the fluid and structure domains, address issues related to the strong and loose coupling schemes, and present a relaxation scheme for the efficient implementation of the strong coupling approach. Finally we present the algorithm for the efficient identification of fluid and solid grid nodes.

For the sake of clarity we adopt the following notation. Lower case variables represent fluid quantities while upper case variables represent structural quantities. Boldface variables are vectors.

2.1. Governing equations for the partitioned FSI problem

The governing equations for the fluid domain are the 3D, unsteady incompressible continuity and Navier–Stokes equations, which in compact tensor notation read as follows:

$$\frac{\partial u_i}{\partial x_i} = 0 \tag{1}$$

$$\frac{\mathrm{d} u_i}{\mathrm{d} t} = -\frac{\partial p}{\partial x_i} + \frac{1}{Re}\frac{\partial^2 u_i}{\partial x_j \partial x_j} \tag{2}$$

where u_i are the Cartesian velocity components, p is the pressure divided by the density ρ, and Re is the Reynolds number of the flow based on a characteristic length and velocity scale. $\mathrm{d}/\mathrm{d}t$ is the material derivative defined as:

$$\frac{\mathrm{d}}{\mathrm{d}t}(\cdot) = \frac{\partial}{\partial t}(\cdot) + u_j\frac{\partial}{\partial x_j}(\cdot)$$

In this paper we consider 3D motions of rigid bodies. For the sake of generality we assume that the rigid bodies are elastically mounted and that the system has damping. Our algorithm is general and can be applied to multiple rigid bodies. For the sake of convenience but without loss of generality, however, we will present the FSI formulation for a single rigid body. Under these assumptions, the Newton's equation of motion that governs the motion of the structural components can be formulated in the inertial frame of reference in generalized form as follows:

$$M\frac{\partial^2 Q_i}{\partial t^2} + C\frac{\partial Q_i}{\partial t} + KQ_i = F_i + F_i^{\text{ext}} \tag{3}$$

where Q_i are the components of the unknown Lagrangian vector $\mathbf{Q}(t)$ describing the position of the structure. For pure translational motion: Q_i is the Cartesian position vector ($Q_i \equiv X_i$) of the rigid body, M is the mass of the object, C is the damping coefficient, K is the spring stiffness coefficient, F_i and F_i^{ext} are the components of the force $\mathbf{F}$ and $\mathbf{F}^{\text{ext}}$ acting on the rigid body from the surrounding fluid and external sources, respectively. For pure rotational motion, $\mathbf{Q}$ is the vector with components the relative angles of the rigid body ($Q_i \equiv \Theta_i$), M represents the moment of inertia, and $\mathbf{F}$ and $\mathbf{F}^{\text{ext}}$ are the moment vectors around the rotation axis arising from fluid-induced and external forces, respectively. For six degree of freedom motion, both the linear and angular Newton's equations need to be solved simultaneously to describe the motion of the rigid body.

The fluid and structure dynamics are coupled together at the fluid/structure interface Γ by the following boundary conditions:

$$u_i = U_i = \frac{\partial X_i}{\partial t} \qquad \text{at} \quad \Gamma \tag{4}$$

Eq. (3) comprises a system of second-order ordinary differential equations. Numerically these equations are typically solved by first transforming them into a system of first-order ordinary differential equations as follows:

$$\frac{\partial Q_i}{\partial t} = \Phi_i \tag{5}$$

$$\frac{\partial \Phi_i}{\partial t} + \frac{C}{M}\Phi_i + \frac{K}{M}Q_i = \frac{F_i + F_i^{\text{ext}}}{M} \tag{6}$$

2.2. Loose and strong coupling of the fluid and structure domains

As we have already discussed, for partitioned FSI solvers either LC-FSI and SC-FSI algorithms can be adopted. The main difference between these two approaches lies in the way the governing equations for the structure domain are integrated in time, i.e., explicit vs. implicit.

To demonstrate this fundamental difference let us revisit Eq. (3). The coupling between the fluid and structure domains arises from the term in the right hand side of this equation **F**. As mentioned above this term is the force or moment exerted by the fluid on the structure and as such it is a function of the position of the structure within the fluid domain, which is determined by **Q** and the flow variables at this location. Recognizing this functional dependence, Eq. (3) can be reformulated as follows:

$$M\frac{\partial^2 Q_i}{\partial t^2} + C\frac{\partial Q_i}{\partial t} + KQ_i = F_i(\mathbf{Q}, \mathbf{q}) + F_i^{\text{ext}} \tag{7}$$

where **q** is the vector containing all flow variables, i.e., $\mathbf{q} = (\mathbf{u} \quad p)^{\text{T}}$. The temporal evaluation of the quasi-linear source term in the right hand side of this equation is essentially what differentiates the LC-FSI and SC-FSI approaches. Consider the following general semi-discrete approximation of Eq. (7):

$$M\frac{\partial^2 Q_i}{\partial t^2} + C\frac{\partial Q_i}{\partial t} + KQ_i = aF_i(\mathbf{Q}^n, \mathbf{q}) + bF_i(\mathbf{Q}^{n+1}, \mathbf{q}) + F_i^{\text{ext}} \tag{8}$$

where a and b are positive constants ($a + b = 1$) determining the temporal accuracy of the scheme and the type of FSI coupling. LC-FSI approaches set $a = 1$ and $b = 0$ while SC-FSI algorithms employ $b \neq 0$. That is, for LC-FSI the fluid-induced force or moment that will determine the position of the structure at time step $n + 1$ is calculated explicitly based on the known location of the structure at time step n. Consequently the structure's location at $n + 1$ can be determined independently of the flow solution at $n + 1$ by solving the now linear Eq. (8). For a SC-FSI implementation, on the other hand, the solution of the structure dynamics and the flow field need to be advanced to time step $n + 1$ in a coupled manner and fluid/structure sub-iterations are required.

The specific iteration schemes we adopt in this work to solve the system given by Eq. (8) for the LC-FSI and SC-FSI algorithms are described in the remainder of this section. Assuming that the solutions for both the fluid $\mathbf{q}^n$ and structure $\mathbf{Q}^n$ domains are known at time step n, the two algorithms determine the solution at the next time step $n + 1$ as follows:

2.2.1. LC-FSI

(1) Using the known position of the structure $\mathbf{Q}^n$ to prescribe boundary conditions for the fluid domain, solve the Navier–Stokes Eq. (2) to obtain $\mathbf{u}^{n+1}$ and p^{n+1}
(2) Calculate the force (or moment) imparted on the solid by the fluid flow $\mathbf{F} = \mathbf{F}(\mathbf{Q}^n, \mathbf{q}^{n+1})$
(3) Solve the system of Eq. (8) using the following temporal integration scheme to obtain $\mathbf{Q}^{n+1}$:

$$\frac{Q_i^{n+1} - Q_i^n}{\Delta t} = \frac{1}{2}(\Phi_i^{n+1} + \Phi_i^n)$$

$$\frac{\Phi_i^{n+1} - \Phi_i^n}{\Delta t} + \frac{C}{M}\Phi_i^{n+1} + \frac{K}{M}Q_i^{n+1} = \frac{F_i + F_i^{\text{ext}}}{M}$$

2.2.2. SC-FSI

(1) Set $\widetilde{\mathbf{Q}}^1 \equiv \mathbf{Q}^n$ and $\widetilde{\mathbf{q}}^1 \equiv \mathbf{q}^n$
   Starting from $\ell = 1$, loop over ℓ(steps a to e) until convergence is achieved:
   (a) Using the known position of the structure $\widetilde{\mathbf{Q}}^\ell$ to prescribe boundary conditions for the fluid domain, solve the Navier–Stokes Eq. (2) to obtain $\tilde{\mathbf{u}}^\ell$ and $\tilde{p}^\ell$.
   (b) Calculate the force on structure domain as $\widetilde{F}^\ell = \mathbf{F}(\widetilde{\mathbf{Q}}^\ell, \widetilde{\mathbf{q}}^\ell)$
   (c) Solve the structure domain equation as

$$\frac{\widetilde{Q}_i^{\ell+1} - Q_i^n}{\Delta t} = \frac{1}{2}(\widetilde{\Phi}_i^{\ell+1} + \Phi_i^n)$$

$$\frac{\widetilde{\Phi}_i^{\ell+1} - \Phi_i^n}{\Delta t} + \frac{C}{M}\widetilde{\Phi}_i^{\ell+1} + \frac{K}{M}\widetilde{Q}_i^{\ell+1} = \frac{\widetilde{F}_i^\ell + \widetilde{F}_i^{\text{ext}}}{M}$$

(d) Check for convergence of the structure solution:

$$||\widetilde{\mathbf{Q}}^{\ell+1} - \widetilde{\mathbf{Q}}^{\ell}|| < \varepsilon$$

where ε is a preset convergence ratio set equal to $\varepsilon = 10^{-6}$ and $||\cdot||$ is the infinity norm in all our simulations.

(e) If not converged, increment ℓ by one and return to step a to continue the iteration loop. If converged, then end the loop and go to 2.

(2) Set $\mathbf{Q}^{n+1} = \widetilde{\mathbf{Q}}^{\ell+1}, \mathbf{u}^{n+1} = \tilde{\mathbf{u}}^{\ell+1}, p^{n+1} = \tilde{p}^{\ell+1}$ and $\mathbf{F}^{n+1} = \widetilde{\mathbf{F}}^{\ell+1}$

It is evident from the above two algorithms that the SC-FSI implementation can be substantially more expensive than LC-FSI. SC-FSI not only requires the iterative solution of the structural equations but within every such iteration the equations governing the fluid motion need to be fully converged. LC-FSI, on the other hand, is very attractive from the computational standpoint since it only requires the solution of the Navier–Stokes equations once per time step. However, loose coupling is known to experience numerical stability and robustness difficulties due to its inherently explicit nature [33]. Therefore, whether one chooses loose or strong coupling FSI iteration is largely dictated by the trade-off between numerical stability and computational cost.

2.3. The added-mass effect

The main issue affecting the stability of LC-FSI algorithms is the so-called "added mass" effect, which models the interaction between the structure and the pressure fluctuations in the fluid [34]. In this section we briefly discuss the added mass issue in order to highlight its importance in the stability of a FSI algorithm and to introduce some concepts that will facilitate a more rigorous stability analysis that will be presented in Section 3.2.3.

The instantaneous force $\mathbf{F}$ imparted on the structure by the fluid arises from the pressure field and the viscous stresses on the fluid/structure interface and can thus be decomposed as follows:

$$\mathbf{F} = \mathbf{F}_{\mathrm{p}} + \mathbf{F}_{\mathrm{v}} \tag{9}$$

where $\mathbf{F}_{\mathrm{p}}$ and $\mathbf{F}_{\mathrm{v}}$ denote the pressure and viscous contributions, respectively. For sufficiently high Reynolds number, the pressure contribution is dominant. The pressure field on the fluid structure interface can be obtained by applying the normal momentum equation, which reads as follows:

$$-\frac{\partial p}{\partial n} = \left[\frac{\mathrm{d}\mathbf{u}}{\mathrm{d}t} - \frac{1}{Re}\nabla^2\mathbf{u}\right] \cdot \mathbf{n} \tag{10}$$

where $\mathbf{n}$ is the unit vector normal to the interface. Using the boundary condition at the fluid/solid interface given by Eq. (4) and noting that at the interface the material derivative of the Eulerian fluid velocity vector $\mathbf{u}$ is the same as the time derivative of the Lagrangian velocity vector of the solid, we can re-write the above equation as follows:

$$-\frac{\partial p}{\partial n} = \frac{\partial^2\mathbf{Q}}{\partial t^2} \cdot \mathbf{n} - \frac{1}{Re}\nabla^2\mathbf{u} \cdot \mathbf{n} \tag{11}$$

By integrating this equation on the interface the pressure distribution on the solid p_s is obtained and the force $\mathbf{F}_{\mathrm{p}}$ can be obtained by integrating the so computed pressure over the interface: $\mathbf{F}_{\mathrm{p}} = \rho \int_{\Gamma}(-p_s\mathrm{d}A)\mathbf{n}$, where $\mathrm{d}A$ is the infinitesimal area of the fluid/solid interface Γ. Therefore, following [34] and using Eq. (9) the resulting force can be formulated as follows:

$$\mathbf{F} = \mathbf{F}_{\mathrm{flow}} - \rho\underline{\mathbf{H}}\frac{\partial^2\mathbf{Q}}{\partial t^2} \tag{12}$$

where $\underline{\mathbf{H}}$ is a linear matrix representing the *added mass* effects and it only depends on the geometry of the structure. As shown in [34] $\underline{\mathbf{H}}$ is a symmetric positive definite matrix. The force $\mathbf{F}_{\mathrm{flow}}$ contains the rest of the fluid force consisting of contributions of the convective and viscous terms to the pressure force as well as the force due to the viscous stresses.

By substituting Eq. (12) into Eq. (3) we can reformulate the dynamic equation for the structure as follows:

$$M\frac{\partial^2 \mathbf{Q}}{\partial t^2} + C\frac{\partial \mathbf{Q}}{\partial t} + K\mathbf{Q} = -\rho(\underline{\mathbf{H}}\frac{\partial^2 \mathbf{Q}}{\partial t^2}) + \mathbf{F}_{\text{flow}} + \mathbf{F}^{\text{ext}} \tag{13}$$

It is straightforward to see from the above equation that when the added mass term in the right hand side is comparable to the natural mass M of the system in the left hand side of Eq. (13), then both the added and natural terms are equal contributors to the dynamics of the structural response. Furthermore, in the limiting situation of systems with very low mass ($M \to 0$), the added terms in the right hand side are the dominant terms that essentially governs the dynamic response of the structure. Recall, however, that the added mass term is embedded in the force (or moment) term $\mathbf{F}$ in the right hand side of the dynamic Eq. (3) and its numerical treatment is essentially dictated by the numerical treatment of $\mathbf{F}$. As mentioned above, LC-FSI evaluates $\mathbf{F}$ explicitly in terms of the known location of the structure from the previous time step. As such, in situations when the added mass becomes of equal or greater importance to the natural terms of the system LC-FSI should be expected to exhibit numerical stability problems. SC-FSI on the other hand handles $\mathbf{F}$ fully implicitly and as such it should be the method of choice for problems with significant added mass effects. As we discuss in the subsequent section, however, even SC-FSI needs to be enhanced with special stabilization and convergence acceleration techniques in order to yield robust and efficient solutions for problems with significant added terms.

The importance of the added mass on the overall stability of LC-FSI has already been shown in FSI of internal flows such as blood flow through compliant vessels [33,35]. In Section 3.1.2 below we will further show the same adverse effects of the added mass on the stability of the LC-FSI approach in FSI of external flows over moving rigid bodies when the rigid body has a very low mass. In such cases, we will demonstrate that strong coupling FSI is essential for obtaining stable and robust solutions. In the stability analysis we report in Section 3.2.3 below we will further show, however, that the relative importance of the structure's mass is not the only parameter that dictates the stability of a FSI algorithm. The local flow properties, which are embedded in the force $\mathbf{F}_{\text{flow}}$ in the right hand side of Eq. (13) can also play a significant role.

2.3.1. Stabilization and convergence acceleration for the SC-FSI

The previously described SC-FSI sub-iterations comprise a Gauss-Seidel like iteration scheme. As shown in [33], for FSI problems involving structures with very low mass (i.e., when the added mass becomes dominant), for which LC-FSI should be expected to be unstable, even the SC-FSI sub-iterations maybe very hard to converge. For such problems Le Tallec and Muro [6] suggested that an under-relaxation scheme for the structure domain is required, which reads as follows:

$$\widetilde{\mathbf{Q}}^{\ell+1} = (1-\alpha)\widetilde{\mathbf{Q}}^{\ell} + \alpha\mathbf{Q}^{\ell+1} \tag{14}$$

In the above equation α is the under-relaxation coefficient ($0 \leqslant \alpha \leqslant 1$), whose value is known to largely determine the overall convergence rate of Gauss-Seidel type iteration schemes. The optimal value of α, however, is problem dependent and for non-linear problems could even vary between different time steps. To remedy this problem, Mok et al. [36] considered two non-linear methods for automatically selecting problem-specific optimal values of α: the gradient method, which chooses the local optimal α searching along a direction based on the solution, and the Aitken acceleration technique [37]. As shown in [36], both methods can drastically improve the convergence rate of the strong coupling FSI iteration relative to sub-iterations with $\alpha = 1$. In this work we choose to implement the Aitken acceleration scheme because of its overall simplicity and computational expedience. The Aitken scheme is implemented as follows:

$$\Delta\widetilde{\mathbf{Q}}^{\ell+1} = \widetilde{\mathbf{Q}}^{\ell} - \mathbf{Q}^{\ell+1} \tag{15}$$

$$\lambda^{\ell+1} = \lambda^{\ell} + (\lambda^{\ell} - 1)\frac{(\Delta\widetilde{\mathbf{Q}}^{\ell} - \Delta\widetilde{\mathbf{Q}}^{\ell+1})^T \Delta\widetilde{\mathbf{Q}}^{\ell+1}}{\|\Delta\widetilde{\mathbf{Q}}^{\ell} - \Delta\widetilde{\mathbf{Q}}^{\ell+1}\|^2} \tag{16}$$

$$\alpha = 1 - \lambda^{\ell+1} \tag{17}$$

For the first iteration ($\ell = 1$), $\Delta\widetilde{\mathbf{Q}}^{\ell+1}$ can be calculated while the $\lambda^{\ell+1}$ cannot be calculated and a preset value of $\lambda = 0.3$ is used to calculate the first under-relaxation coefficient α.

2.4. The flow solver

In this section we present the numerical method for solving the Navier–Stokes equations in a domain with multiple moving immersed boundaries. The basic method is the recently published CURVIB approach of Ge and Sotiropoulos [2], which is only briefly described below. Most emphasis in this section is placed on the description of an efficient and general algorithm for classifying the nodes of the background mesh according to their position in the fluid or structure domains.

2.4.1. The CURVIB method

The Navier–Stokes and continuity equations, Eq. (2), in a domain containing multiple, 3D, arbitrarily complex moving immersed boundaries are solved using the sharp-interface CURVIB (curvilinear grid/immersed boundary) solver of Ge and Sotiropoulos [2]. In the CURVIB approach a curvilinear grid system is adopted to serve as the background grid and the immersed bodies are treated as sharp-interface immersed boundaries. The curvilinear background mesh is adopted to enhance algorithmic flexibility and efficiency for internal flow problems in which the background domain can be efficiently discretized with a boundary-conforming curvilinear mesh (see [2] for details).

The method employs a novel, fully-curvilinear staggered grid discretization approach, which does not require either the explicit evaluation of the Christoffel symbols or the discretization of all three momentum equations at cell interfaces as done in previous curvilinear staggered grid formulations. The equations are integrated in time using an efficient, second-order accurate fractional step methodology coupled with a Jacobian-free, Newton–Krylov solver for the momentum equations and a GMRES solver enhanced with multigrid as preconditioner for the Poisson equation. For more details about the fluid solver the readers are referred to [2].

2.4.2. An efficient grid node classification algorithm

The key feature of the CURVIB approach, and for that matter of any sharp-interface immersed boundary method, is that the governing equations are solved on a grid system that does not conform with the boundaries of the immersed bodies (see Fig. 2). The governing equations are solved only at nodes located within the fluid (fluid nodes) with boundary conditions reconstructed at fluid nodes that are in the immediate vicinity of the immersed boundary (immersed boundary (IB) nodes). The background grid nodes that are located within the solid (solid nodes) are blanked out of the computation. Therefore, the first step in the implementation of the CURVIB approach is the classification of the grid nodes of the background mesh into these three categories: fluid, IB, and solid nodes. For FSI problems involving complex 3D bodies this step could be critical for the overall efficiency of the computational algorithm as it needs to be repeated each time the location of the structure within the flow domain is updated. In what follows we describe an efficient and general algorithm for classifying grid nodes.

The task of identifying whether a grid node is fluid or solid node is the classical problem of point-in-polyhedron problem in computational geometry: given a point P in space and a polyhedron, whose geometry is defined by it boundaries, find out whether the point is located inside or outside the polyhedron. For 2D problems, the boundaries of objects are described by polygons and the problem is downscaled to 2D
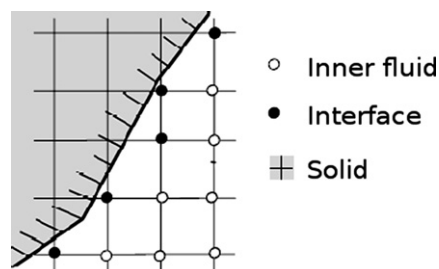

Fig. 2. Sketch showing the idea of sharp-interface immersed boundary method.

point-in-polygon problem. Several methods have been developed in the past to solve the point-in-polygon problem and the most widely used among them are the so-called ray-casting method and the winding number method [32]. The winding number method usually involves the solution of inverse trigonometric functions and thus it is not as efficient as the ray-casting method. The ray-casting method is based on the Jordan curve theorem [38] and can be summarized as follows. Starting from the point under consideration, cast a random half-infinite ray and count the number of intersections between the half ray and polygon edges. If the number of intersections is odd then the point is located inside the polygon (point A in Fig. 3), otherwise it is located outside (point B in Fig. 3). Extending this idea to 3D leads to the following point-in-polyhedron algorithm.

Let the surface of the polyhedron be discretized with an unstructured mesh consisting of a set of triangles, Δ_κ, $\kappa = 1, K_{max}$ where K_{max} is the total number of triangles. The point to be identified is P with Cartesian coordinates (P_1, P_2, P_3). The point-in-polyhedron identification can be carried out by casting a line from P to a point S, which is located outside the polyhedron, and count the number of intersections. The algorithm is described in the following pseudo-code:

Algorithm 1. Point-in-Polyhedron test
    $count = 0$
    for $\kappa = 0$ to K_{max} **do**
        if PS intersect with triangle Δ_κ **then**
            $count = count + 1$
        end if
    end for
    if $count$ is odd **then**
        P is a solid node
    else
        P is a fluid node
    end if

The core element in the above algorithm is the ray-triangle intersection testing algorithm. In this work we adopt the ray-triangle intersection method of [39] (see Fig. 4). The basic idea of the method is that for a triangle defined by its three vertices $V^j = (v_1^j, v_2^j, v_3^j)(j = 1, 2, 3)$, where and v_i^j is the ith Cartesian coordinate of the jth vertex, any point $T = (\tau_1, \tau_2, \tau_3)(i = 1, 2, 3)$, where τ_i is the ith Cartesian component of point T, within the triangle can be expressed as follows:

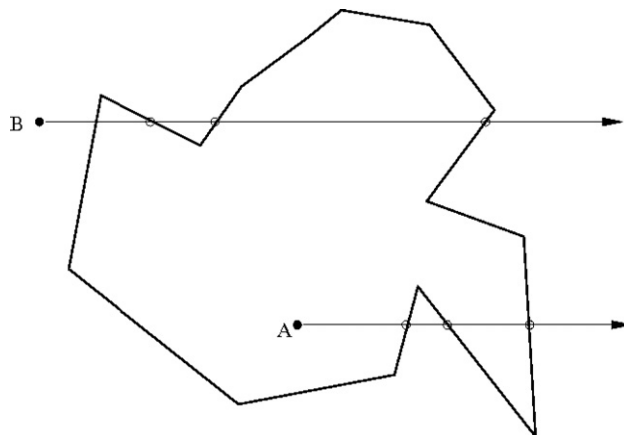$$\tau_i(a, b) = (1 - a - b)v_i^1 + av_i^2 + bv_i^3 \tag{18}$$

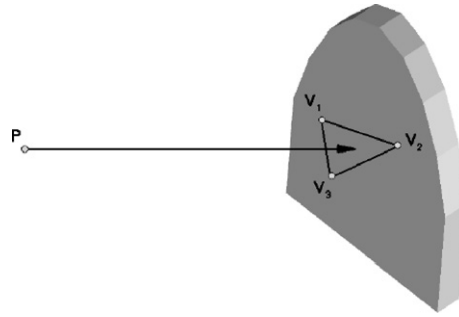Fig. 3. Sketch showing the ray-casting method for 2D point-in-polygon test.

Fig. 4. Sketch showing the ray-casting method for a 3D leaflet.

with $a \geqslant 0$, $b \geqslant 0$ and $a + b \leqslant 1$. The ray $\mathbf{PS} = (PS_1, PS_2, PS_3)$, on the other hand, can be expressed as follows:

$$PS_i(\delta) = P_i + \delta e_i \tag{19}$$

where e_i are the components of the unit vector emanating from P and pointing along vector $\mathbf{PS}$ from P to S and $\delta > 0$. Finding the intersection between the ray and a triangle, therefore, is equivalent to finding the solution to the following parametric equation:

$$PS_i(\delta) = \tau_i(a, b) \tag{20}$$

If the solution to Eq. (20) satisfies $\delta > 0, a \geqslant 0, b \geqslant 0$ and $a + b \leqslant 1$ then the ray intersects with a triangle.

The point-in-polyhedron algorithm described above is fairly straightforward to implement but it is quite inefficient. This is primarily due to the fact that the straightforward implementation of the above scheme leads to an algorithm with computational cost that is $O(N)$. To remedy this situation we develop a two step procedure that drastically enhances the efficiency of the point-in-polyhedron algorithm. We illustrate this procedure for the bileaflet mechanical heart valve case, which we investigate in detail in Section 3.2, as follows.

(1) Bounding box
    The bounding box (Fig. 5a) is defined as the smallest Cartesian box that contains all vertices defining the solid object's surfaces, i.e., the size of the bounding box is defined as

    $$[x_{\min}, x_{\max}] \times [y_{\min}, y_{\max}] \times [z_{\min}, z_{\max}] \tag{21}$$

    where $x_{\min} = \min_{i=1,K_{\max}}(x_i)$ and $x_{\max} = \max_{i=1,K_{\max}}(x_i)$. Any grid node located outside the bounding box is outside of the body and consequently it is a fluid node. Depending on the specific problem, this step can potentially reduce the number of point-in-polyhedron tests considerably if a large number of the grid nodes are located outside the bounding box.

(2) Control cell method
    The control cell method is an extension of the "grid cell method" for 2D point-in-polygon detection as described in [32] to 3D. The idea is to cut the bounding box into $I_x \times I_y \times I_z$ uniform control cells, as shown in Fig. 5b, with each control cell having the dimension of $\Delta x = L_x/(I_x - 1)$, $\Delta y = L_y/(I_y - 1)$, and $\Delta z = L_z/(I_z - 1)$, where L_x, L_y, and L_z depicts the length of the bounding box along x-, y- and z-direction, respectively. Each control cell is categorized as either an empty cell, which does not intercept with any surface triangles, or a full cell, which intersects with at least one of the surface triangles. Each empty control cell can be further defined as fluid cell (entirely outside of the polyhedron) or solid cell (inside the polyhedron). For each full cell an array containing the complete list of all intersecting triangles is created. In the standard ray-casting algorithm for the point-in-polyhedron test, the ray is allowed to be shot in an arbitrary manner and consequently we need to screen through all triangles to find out all possible intersections between the ray and surface triangles. Following the creation of the control cells,
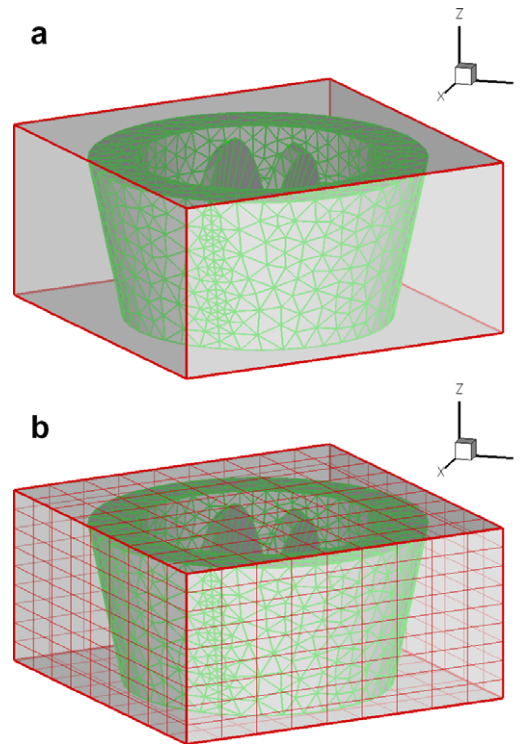
Fig. 5. (a) A bounding box created for a bileaflet mechanical valve. The valve is discretized with 2048 triangles. (b) The bounding box is divided into smaller control cells.

however, we can fix the ray shoot process toward a predefined direction. The reason for doing this is that by fixing the ray shooting direction we know that this ray is only going to interact with a limited number of control cells that can be easily identified. For example, without lose generality, we choose the positive z direction as the ray shooting direction. Assume that the point P is located within control cell $(1, 1, 1)$. The line segment PS is thus only possible to interact with all control cells from $(1, 1, 1)$ to $(1, 1, I_z)$ and for that we only need to carry out the ray-triangle intersection test for those triangles that intersect with these control cells. The complete point-in-polyhedron algorithm we adopt in this work can be described as follows:

It is easy to show that, assuming that the surface triangles are randomly and uniformly distributed in space, for each search point P the computational cost of this point-in-polyhedron scheme is about $O(N/(I_x \times I_y))$, which could be orders of magnitude lower than the original $O(N)$ scheme. For the mechanical heart valve problem we consider below in Section 3.2, we found that choosing $I_x = I_y = I_z = 20$ is sufficient to yield very fast grid node identification speed. More specifically, for a background mesh consisting of 10 M grid nodes and with the leaflets of the heart valve discretized with 2048 triangles the cost of the grid node classification step using the current identification algorithm is less than 5% of the total FSI algorithm computational cost per time step.

3. Numerical results

In this section we apply the numerical method to simulate two FSI problems. The first case is the vortex induced vibration of elastically mounted cylinders and the second case is the FSI problem of blood flow through a bileaflet mechanical heart valve. These cases are selected here to: (1) validate the FSI solver; (2) illustrate the capability of the FSI solver in handling complicated FSI problems; and (3) investigate the estability and robustness of the FSI solver.

Algorithm 2. Modified Point-in-Polyhedron test

  if P is within the bounding box **then**

    Locate the control cell (i_x,i_y,i_z) where point P is l0ocated with

$$i_x = INT((x_p - x_{\min})/\mathrm{d}x) + 1$$
$$i_y = INT((y_p - y_{\min})/\mathrm{d}y) + 1$$
$$i_z = INT((z_p - z_{\min})/\mathrm{d}z) + 1$$

    if cell (i_x,i_y,i_z) is an empty cell **then**

      P is a fluid node if cell (i_x,i_y,i_z) is a fluid cell, otherwise it is a solid node

    else

      count $= 0$

      for $k = i_z$ to I_Z **do**

        for each triangle Δ intercepting with the control cell (i_x,i_y,k) **do**

          if PS intercepts with Δ **then**

            count $=$ *count* $+ 1$

          end if

        end for

      end for

      if *count* is odd **then**

        P is a solid node

      else

        P is a fluid node

      end if

    end if

  end if

3.1. Vortex-induced vibrations (VIV) of 2D elastically mounted cylinders

Vortex-induced vibrations (VIV) of elastically mounted obstacles are of interest in a broad range of areas of engineering practice. These include, among others, the design of risers and conductor tubes in oil drilling platforms, civil structures such as bridges and chimneystacks, marine structures in the ocean, etc. In addition to their great practical relevance VIV problems are also important from a fundamental fluid mechanical standpoint due to the enormous richness of their underlying vorticity dynamics. For these reasons a large and continuously expanding body of literature is dedicated to the study of VIV–see the recent review by [40].

The flow past a single, elastically mounted 2D cylinder has served as the generic VIV model problem and has been widely studied both numerically and experimentally. As shown in Fig. 6, a 2D circular cylinder with one-degree of freedom is elastically mounted in a uniform flow of velocity U, with the cylinder's free vibration direction perpendicular to the flow direction. The cylinder diameter is D. The Reynolds number of the flow is defined as $Re = UD/v$ where v is the kinematic viscosity of the fluid. The dynamic equation governing the motion of the cylinder reads in dimensional form as follows:

$$M\frac{\partial^2 Y}{\partial t^2} + C\frac{\partial Y}{\partial t} + KY = F_Y \tag{22}$$

where Y is the coordinate of the cylinder center along y-axis, M is the mass, C damping factor, K is stiffness of the spring and F_y is the total fluid force on the cylinder in the Y direction. The structure's natural frequency and critical damping factor are given by the following equations:

$$\omega = 2\pi f = \sqrt{\frac{K}{M}} \tag{23}$$
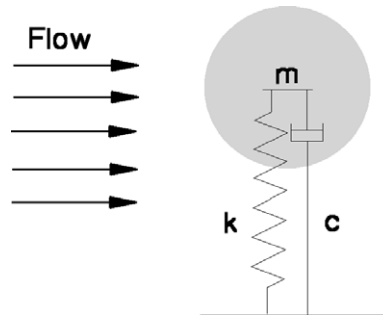$$C_{cr} = 2\sqrt{MK} = 2K\omega \tag{24}$$

Fig. 6. An elastically mounted cylinder in the free stream.

By appropriately grouping its coefficients into non-dimensional groups, Eq. (22) can be further formulated as follows:

$$\frac{\partial^2 Y}{\partial t^2} + 4\pi\xi \frac{1}{U_{\text{red}}} \frac{\partial Y}{\partial t} + 4\pi^2 \frac{1}{U_{\text{red}}^2} Y = \frac{1}{2M_{\text{red}}} C_Y \tag{25}$$

The various non-dimensional coefficients in the above equation are defined as follows:

Non-dimensional damping coefficient

$$\xi = \frac{C}{C_{\text{cr}}} \tag{26}$$

Reduced velocity

$$U_{\text{red}} = \frac{U}{fD} \tag{27}$$

Reduced mass

$$M_{\text{red}} = \frac{M}{\rho D^2} \tag{28}$$

Non-dimensional force coefficient

$$C_Y = \frac{2F_Y}{\rho U^2 D} \tag{29}$$

The solution of the FSI problem for this case is governed by these four non-dimensional numbers and the Reynolds number of the flow. In this work we restrict our attention to a system with no damping and set $\xi = 0$.

3.1.1. Validation

It is well known that when the natural frequency of the cylinder falls within the so-called "lock-in" region, large amplitude vibrations are excited [41]. Within this region, the vortex shedding frequency changes to match the frequency of the structure's motion resulting to the observed large vibration amplitudes. The synchronization frequency is not necessarily the natural frequency of the structure and has often been observed to exceed it significantly [42–44]. To validate the FSI solver we apply it to investigate the lock-in phenomenon for a case for which benchmark numerical results are available in the literature.

The computational domain is a rectangular box with dimensions $32D \times 16D$. The cylinder is initially located on the horizontal axis of symmetry of this box $8D$ from the inlet (as shown in Fig. 7). Unifrom flow is prescribed at the inlet and slip walls are used for the side boundaries. The convective boundary condition is used at the outlet [45]. The computational domain is discretized with 281×241 grid nodes, which are clustered toward the four sides of a square box centered around the cylinder (see Fig. 7). This square box is discretized with a uniform 50×50 mesh, which corresponds to a near-cylinder grid spacing of $\Delta h \sim 0.02D$. A time step of $\Delta t = 0.02$ is used.
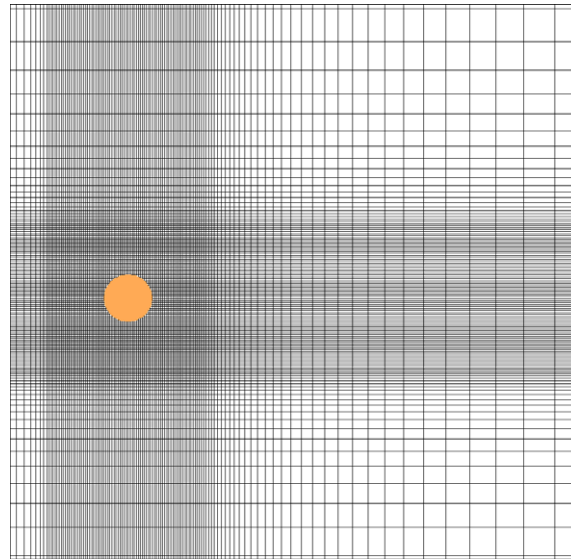
Fig. 7. Grid used for 2D vortex induced vibration (VIV) study around an elastically mounted circular cylinder.

To investigate the "lock-in" phenomenon, we fix the Reynolds number and reduced mass of the system ($Re = 150$ and $M_{red} = 2$) and systematically vary the natural frequency of the system. This is accomplished by varying the reduced velocity U_{red} with increments of 1 within the range of $3 \leqslant U_{red} \leqslant 8$. The resulting variation of the maximum displacement of the cylinder with U_{red} conditions is plotted in Fig. 8. As seen in the figure, large amplitude vibration (exceeding *10* percent of the cylinder diameter) is observed within $U_{red} \in [4, 7]$ while outside this region the cylinder vibration amplitude is drastically reduced. Based on these results the lock-in region for this system is $U_{red} \in [4, 7]$ with maximum vibration amplitude of about $0.5D$ attained for $U_{red} = 4$ (see Fig. 10 for the calculated wake patterns for this case). To quantitatively validate our method, we also include in this figure the recent results of Ahn and Callinderis [28] who employed an unstructured, finite-element ALE approach. As seen in the figure the two results are in excellent agreement with each other.
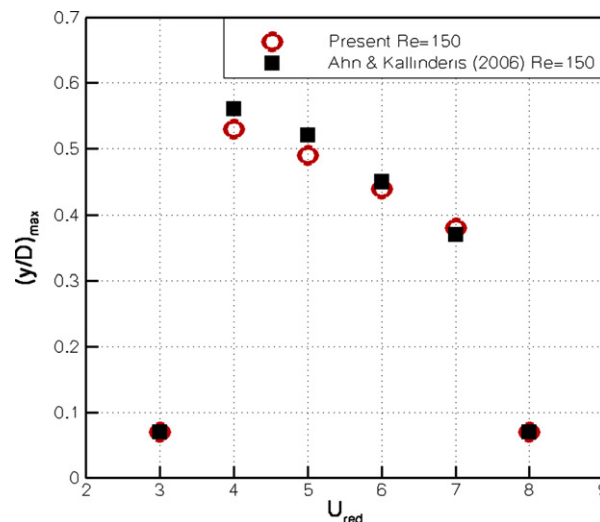

Fig. 8. The maximum cylinder displacement as a function of U_{red} from present (circles) and Ahn and Kallinderis [28] (filled squares) simulations. Flow conditions: $Re = 150$, $M_{red} = 2$.
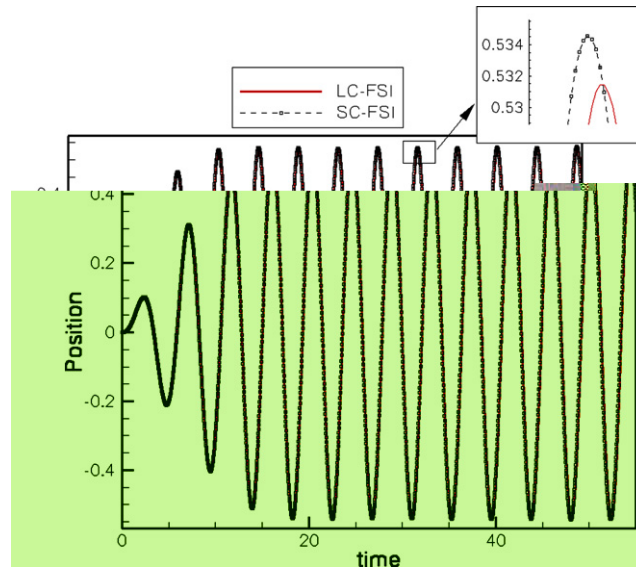
Fig. 9. Comparison of the cylinder kinematics obtained by LC-FSI and SC-FSI. Flow conditions: $Re = 150$, $U_{red} = 4$, $M_{red} = 2$.
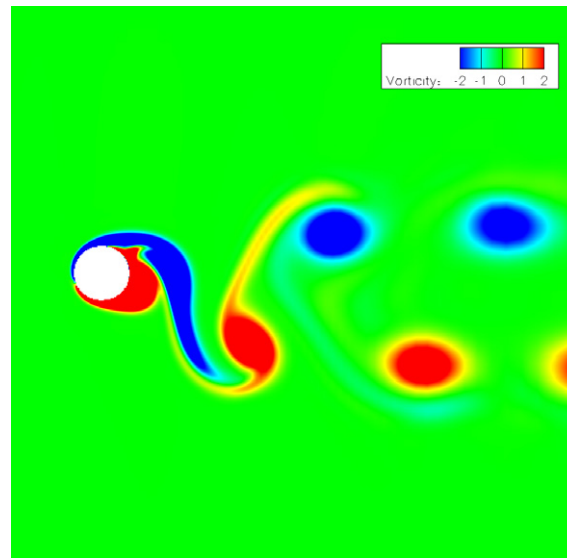

Fig. 10. Instantaneous vorticity contours in the vicinity of the cylinder. Flow conditions: $Re = 150$, $U_{red} = 4$, $M_{red} = 2$.

3.1.2. LC-FSI vs SC-FSI algorithms

The above computations for the entire range of simulated reduced velocities were carried out using the LC-FSI algorithm, which yielded robust and stable solutions for all cases. To examine the importance of the FSI coupling algorithm, we carry out simulations both with strong and loose coupling schemes for $Re = 150$, $M_{red} = 2$, and $U_{red} = 4$ (i.e., the set of parameters for which the maximum vibration amplitude is observed). For the SC-FSI iteration, the under-relaxation scheme of Eq. (14) was not required. The SC-FSI sub-iterations remain robust without any under-relaxation and convergence is reached within 3–4 sub-iterations per time step. The time history of the cylinder center location calculated by the two different coupling approaches is shown in Fig. 9. As seen in the figure, both coupling approaches lead to solutions that are practically indistinguishable from each other. Only minute differences are observed when zooming to very small scales in this plot. We carried out similar comparisons using different time steps and obtained identical results.

The above results concerning the performance of LC-FSI and SC-FSI iterations should not be surprising since in Section 2.3 we argued that the differences between the two algorithms should start appearing only when the reduced mass of the structure becomes sufficiently small. To demonstrate the validity of this hypothesis, we carry out a numerical experiment for an elastically mounted cylinder with the following set of parameters: $Re = 150$, $M_{red} = 0.02$, and $U_{red} = 4$. That is, we reduce the mass of the system by two orders of magnitude. All other numerical parameters (computational grid and time increment) are kept the same as for the $M_{red} = 2$ case. Our results show that the LC-FSI iteration becomes rapidly unstable immediately after the start of the simulation. The same trend is also obtained for the SC-FSI algorithm when no under-relaxation is used. For both algorithms reducing the time step by as much as one order of magnitude appeared to have no effect on the stability of the FSI iteration.

Converged solutions could be obtained only when the SC-FSI algorithm was coupled with the Aitken under-relaxation scheme of Section 2.3.1, typically within 5 FSI sub-iterations per time step. This finding demonstrates the ability of the SC-FSI algorithm with the Aitken under-relaxation strategy to yield converged solutions for systems with very low mass, which can not be obtained otherwise. The significance of this finding will be appreciated and discussed in depth in a subsequent section of this paper where we report results for the prosthetic bileaflet heart valve case – a system that is inherently characterized by the very low moment of inertia of its structural components.

3.1.3. VIV of six elastically mounted 2D cylinders

In this section we carry out a simulation aimed at demonstrating the ability of our method to simulate FSI problems involving many immersed bodies. We carry out a simulation for six elastically mounted cylinders initially placed perpendicular to the flow direction. Each cylinder is allowed to have two degrees of freedom by attaching to two strings in the x and y directions, respectively. The parameters of the system are set as $Re = 100$, $M_{red} = 2$ and $U_{red} = 4$. The interaction between the flow and the cylinders as well as the strong jets forming between the cylinders lead to a very rich wake flow as illustrated in the two snapshots of the instantaneous vorticity field shown in Fig. 11. Both LC-FSI and SC-FSI iterations were found to work successfully for this case and yielded robust and stable solutions for the entire simulated interval.

3.2. FSI simulation of a bileaflet mechanical heart valve at physiologic conditions

3.2.1. Background

The second problem we consider to demonstrate the capabilities of the FSI solver is 3D pulsatile flow through a bileaflet mechanical heart valve (BMHV) implanted in a modeled straight aorta at physiologic flow conditions. The BMHV is today the most widely implanted worldwide prosthesis for replacing diseased natural heart valves because it combines satisfactory overall hemodynamics performance with very long durability [46]. However, all current BMHV designs are prone to a number of complications, including among others increased likelihood of blood clotting and possible hemorrhage [47]. Such complications are believed to be strongly associated with the complex, non-physiologic blood flow patterns induced by the valve.

A typical BMHV, as shown in Fig. 1, consists of a circular housing and two leaflets, which are attached to the housing through a hinge mechanism. When implanted in the aortic position of the human heart, the valve is subjected to a periodic pulsatile flow resulting from the periodic contraction and expansion of left ventricle. For a typical adult patient the Reynolds number of the flow (based on the aorta diameter and the bulk flow at the inflow) varies from zero, during the diastolic phase when the valve is closed, to nearly 6000, at peak systole when the valve leaflets are fully open. The two leaflets, which are free to pivot around their respective hinge mechanism, open and close passively in response to the forces imparted by the ambient flow giving rise to a very complex FSI problem. The complexity of the problem stems both from the geometrical complexity of the entire configuration and the richness of the flow that emerges as the incoming flow pulse interacts with this complex and dynamically evolving geometry.

The first comprehensive insights into the complexity of BMHV hemodynamics was reported in the recent experimental and computational study of Dasi et al. [1], who employed particle image velocimetry (PIV) measurements and direct numerical simulations (DNS) to investigate the flow through a clinical quality BMHV (St. Jude Regent 23 mm valve) implanted in a model, axisymmetric aorta geometry at physiologic, pulsatile
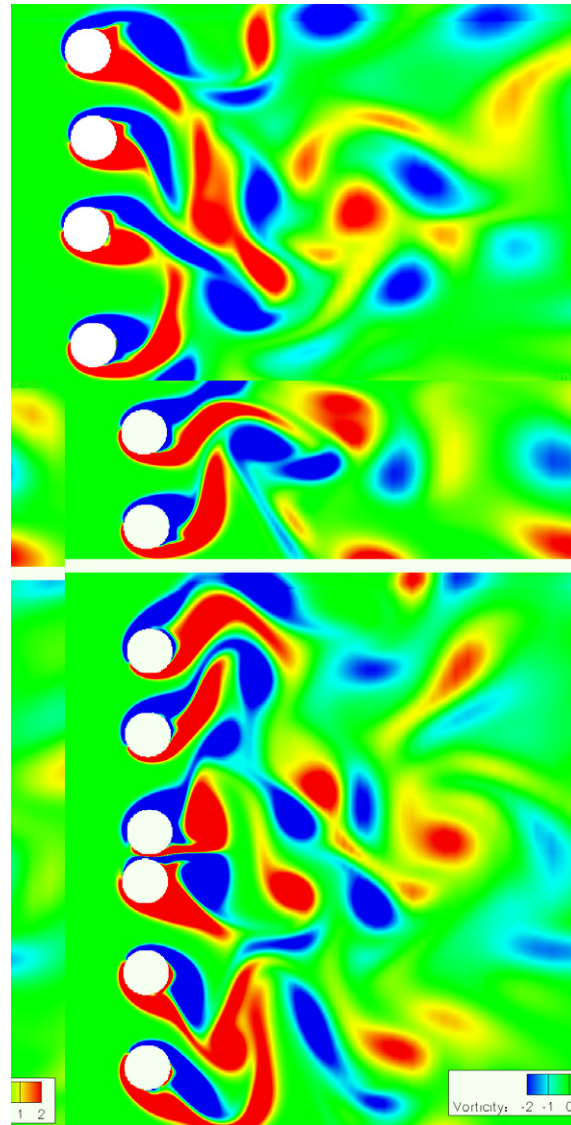
Fig. 11. VIV of six elastically mounted cylinders. Two snapshots of out-of-plane vorticity contours. Flow conditions: $Re = 100$, $U_{red} = 4$, $M_{red} = 2$.

flow conditions (see Fig. 1 for a sketch of the model geometry). The main features of BMHV flows as emerge from the results of Dasi et al. [1] can be summarized as follows: (a) during the early systolic phase when the valve leaflets open rapidly the flow is laminar but is dominated by the emergence of complex 3D vortical structures, which are shed from the valve leaflets and the step-like expansion of the aortic wall into the model axisymmetric sinus; (b) as the flow continues to accelerate, vortex-to-vortex interactions and stretching of vorticity by induced pressure gradients give rise to a very complex, albeit still laminar, flow state. Throughout the accelerating systolic phase, the flow remains well-organized and repeatable from cycle to cycle; (c) very close to peak systole and shortly before the decelerating systolic phase commences, the flow in the wake of the leaflets bursts rapidly into a chaotic, turbulent-like state, which is characterized by large cycle-to-cycle variations; (d) this chaotic state is intensified throughout the flow deceleration and persists, albeit at diminishing intensity, even as the flow enters the diastolic phase and the valve leaflets close; and (e) the remnants of this chaotic flow state are rapidly swept downstream into the aorta as the systolic phase commences again and the valve leaflets open. As such, each new cardiac cycle amounts to a "fresh" start for the flow without residual

disturbances from the previous cycle. All these features of the flow observed in the experiment were captured with good accuracy by the DNS [1]. It is important to emphasize, however, that the DNS reported by Dasi et al. [1] was carried out using the CURVIB approach of [2] by prescribing both the pulsatile inflow conditions and the kinematics of the leaflets from experimental measurements (i.e., the DNS reported in [1] did not tackle the FSI aspects of the problem). We should also note that in [2] we reported similar simulations (i.e., by prescribing the leaflet kinematics) for a BMHV mounted in a straight aorta with a non-axisymmetric, triple-sinus geometry in order to demonstrate the ability of the CURVIB method to yield efficient solutions in complex curvilinear domains with moving immersed boundaries. The experiments of Dasi et al. [1], however, revealed a number of additional important features pertaining to the FSI aspects of the problem, which could not be explored by a DNS with prescribed kinematics. These include the following findings: (a) the valve leaflets open and close very rapidly, typically about 80 ms and 60 ms, respectively; (b) the two leaflets open and close asymmetrically (i.e., one opens or closes faster than the other); and (c) the precise leaflet kinematics can have a major impact on the structure of the wake flow. Exploring these issues computationally requires the use of FSI algorithms. Several studies have been reported in recent years aimed at exploring BMHV flows including FSI phenomena (see [1] for a detailed review) but to the best of our knowledge methods that can carry out highly resolved, FSI simulations of BMHV flows at physiologic flow conditions have yet to be developed. In several recent numerical works the problem has been simplified by considering either 2D geometry [48–51] or 3D geometry with the added assumption of quadrant symmetry [52] (i.e., modeling only one quarter of the valve cross-section and assuming that the flow in the wake of the leaflet will comply with the quadrant symmetry of the geometry). Full 3D simulations have also been reported but have focused on lower Reynolds numbers (well outside the physiologic range) and used coarse numerical resolution [53]. The results we present in this section constitute to the best of our knowledge the first high-resolution, 3D, FSI simulation of BMHV flows at physiologic Reynolds numbers.

3.2.2. Computational details and results

We carry out FSI simulation of the same BMHV problem we investigated in [1]. The valve is placed in a straight model aorta with an axisymmetric sinus expansion as shown in Fig. 1. As in [1], we neglect the hinge mechanism connecting the leaflets and the valve housing. In addition, the damping due to the friction of the hinge has been neglected as well since it is very small relative to the flow forces and moreover no experimental data is available on the actual value of the friction coefficient. Each leaflet is considered to rotate around a fixed hinge axis i.e., one degree of freedom, as illustrated in Fig. 1. The hinge axis in our simulations has been placed parallel to the x-axis of the global frame at y_h and z_h from the origin (see Fig. 1). The motion of the leaflets is governed by the angular momentum equation around the x-axis, which after non-dimensionalizing yields as follows:

$$I_{\text{red}} \frac{\partial^2 \theta}{\partial t^2} = C_{\Im} \tag{30}$$

In the above equation, $\theta = \theta(t)$ is the leaflet angle defined as in Fig. 1, which can vary between $\theta_{\min} = 5°$ and $\theta_{\max} = 54°$. I_{red} is the leaflet reduced moment of inertia defined as:

$$I_{\text{red}} = \frac{\rho_s I_{xx}}{\rho_f D^5} \tag{31}$$

where, D is diameter of the aorta, ρ_f is the density of the fluid, ρ_s is the density of the leaflet material, and $I_{xx} = \int r^2 \, dV$ (with $r = \sqrt{(y - y_h)^2 + (z - z_h)^2}$ and dV as the infinitesimal volume element) is the product of inertia around the leaflet's hinge axis. The reduced moment of inertia is calculated to be $I_{\text{red}} = 0.001$, by assuming Polycarbonate as the leaflet material with density of 1750 kg/m^3 and water as fluid with the density of 1000 kg/m^3. The moment of inertia I_{xx} is approximated by assuming that the leaflet is a short half cylinder – the leaflets are actually manufactured using half cylinders by trimming off near the hinge and the trailing edge regions. The moment coefficient exerted by the flow on the leaflet is defined as

$$C_{\Im} = \frac{\Im}{\rho_f U^2 D^3} \tag{32}$$

where, $\Im$ is the moment around the hinge axis and U is the bulk flow velocity.

From the FSI simulation standpoint and based on our previous discussion in Section 3.1.2 and the results we reported above for the 2D cylinder problem, the very low moment of inertia $I_{red} = 0.001$ of the BMHV leaflets should be expected to pose a major computational challenge. The results we will present below confirm that this is indeed the case.

To discretize the empty aorta domain we employ two body-fitted curvilinear grid systems: a fine mesh with $201 \times 201 \times 241 \approx 9.7 \times 10^6$ nodes; and a coarse mesh with $153 \times 153 \times 241 \approx 5.6 \times 10^6$ grid nodes (see Fig. 12). The two grids are used to evaluate the sensitivity of the computed solutions to mesh refinement as described later in this section. Unless indicated otherwise in our subsequent discussion, however, most results reported below have been obtained on the fine mesh. For both grids the valve leaflets and housing are discretized with a triangular mesh as needed by the node classification algorithm (see Fig. 12). The valve is about $4D$ from the inlet and total domain length is $15D$. At the outflow boundary, the convective boundary condition is used [45]. At the inflow boundary we prescribe unsteady, pulsatile plug flow based on the experimental data reported in [1]. Fig. 13 shows the corresponding experimental time history of flow rate within a cardiac cycle, which we prescribe as inflow condition. According to the experiment [1], the time period of the cardiac cycle is 860 ms and the peak Reynolds number based on the upstream pipe diameter and peak bulk velocity is 6000. The cardiac cycle can be divided into two phases: a systolic phase ($t = 0$–350 ms) and a diastolic phase. At the onset of systolic phase, the positive incoming flow opens the valve, while near the transition from systolic to diastolic phase the flow direction changes for a brief time interval thus driving the two leaflets toward to closing position. A physical time step of $\Delta t = 0.33$ ms, corresponding to about 2580 time steps per cardiac cycle, is selected to advance the simulation in time. All simulations start from a zero flow initial condition with the valve leaflets in the closed position and the prescribed inflow is accelerated according to the wave form shown in Fig. 13. Simulations are carried out both with the LC-FSI and SC-FSI algorithms. Our findings concerning the stability of the two algorithms are to some extent consistent with the results we obtained for the low reduced mass cylinder case reported above. However, we also find some unexpected
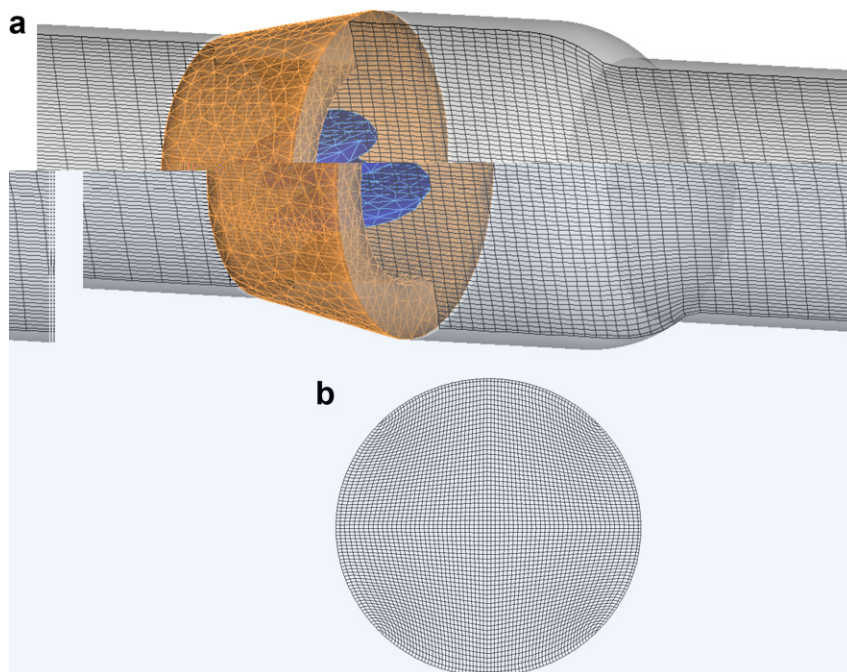

Fig. 12. The computational grid for the BMHV simulations. (a) Side view showing the background curvilinear grid used to discretize the empty aorta and the unstructured mesh used to discretize the leaflets and the housing. (b) Cross-sectional view of the background grid. For clarity, only every third grid is plotted.
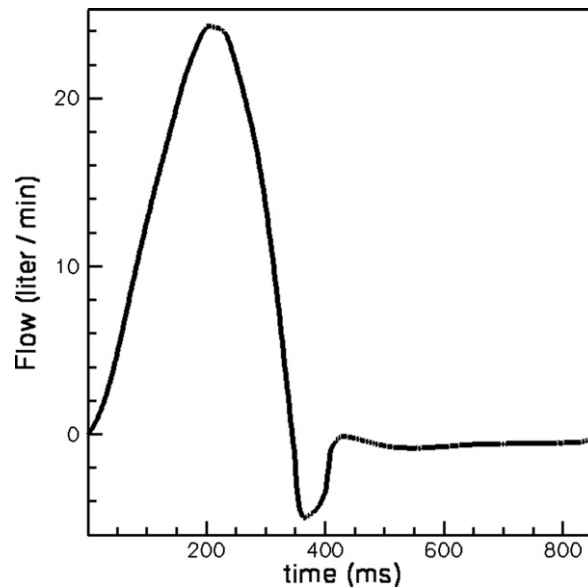
Fig. 13. BMHV simulation. Physiological incoming flow waveform specified at the inlet.

trends, which underscore the complexities of carrying out FSI simulations of 3D problems involving complex flow phenomena.

We initially attempted to solve the problem using the LC-FSI approach but the computation blew up within few simulated time steps of the valve opening process. Reducing the time step by one order of magnitude to 0.03 ms did not alter the stability of the LC-FSI iteration and the computation also blew up during the very early stages of the opening phase. As we will show in Section 3.2.3 in some cases no matter how small the time step the opening phase computation with LC-FSI is going to be unstable. These findings are entirely expected based on our previous discussion on the detrimental effects of low system mass on the stability of the LC-FSI iteration and the very low inertia of the valve leaflets.

Following the failure of LC-FSI, we attempted the same simulation using the SC-FSI scheme, starting from the same initial conditions, and using the same time step ($\Delta t = 0.33$ ms). Without the use of under-relaxation, Eq. (14), the stability of the SC-FSI sub-iterations were erratic often failing to reach full convergence. The effect of under-relaxation on the performance of the SC-FSI scheme is illustrated in Fig. 14, which shows the convergence histories for various under-relaxation coefficients during one time step in the opening phase of the valve. The results obtained using the Aitken's convergence acceleration technique in Section 2.3.1 are also included in this figure. When constant under-relaxation coefficient α is employed, the convergence rate depends on the magnitude of α. Values close to one (see results for $\alpha = 0.7$ in the figure) yield slow convergence rates (over 20 sub-iterations are required) and the final solution is approached in a non-monotonic manner via oscillations about the "exact" solution. The amplitude of the oscillations is initially very large and it is damped very gradually by the SC-FSI sub-iterations. Reducing the under-relaxation coefficient to $\alpha = 0.5$ clearly enhances the convergence rate of the SC-FSI sub-iterations. Convergence is now reached within 8–9 sub-iterations and the solution is also approached in an oscillatory manner. The implementation of Aitken's acceleration has a profound effect on the performance of SC-FSI yielding monotonic convergence within 5 sub-iterations per time step (similar convergence rates are obtained throughout the simulated cardiac cycle). These results demonstrate the impact of under-relaxation on the efficiency and robustness of SC-FSI and underscore the substantial efficiency gains achieved via the Aitken's acceleration technique. Recall that a major advantage of the Aitken's acceleration is that the method automatically selects the optimal under-relaxation coefficient at every sub-iteration.

The SC-FSI algorithm with the Aitken's acceleration technique remains stable throughout the entire cardiac cycle and consistently yields converged solutions within 4–5 sub-iterations. The so computed leaflet
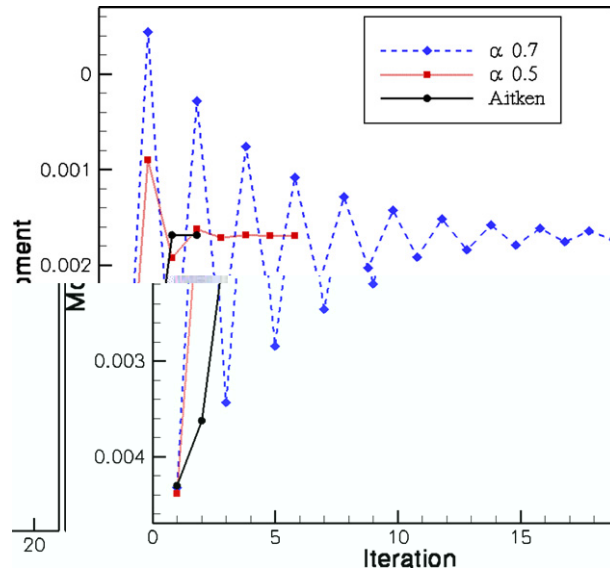
Fig. 14. BMHV simulation, Comparison of the SC-FSI convergence history of calculated moment acting on a leaflet with constant coefficient under-relaxation and Aitken's acceleration within a time step.

kinematics are compared with the experimental measurements of Dasi et al. [1] in Fig. 15. It is seen that our simulations reproduce the measured kinematics with remarkable accuracy both in terms of the overall duration and acceleration rates of the opening and closing phases. A few snapshots of the calculated flowfields are plotted in Fig. 16 in terms of the instantaneous, out-of-plane vorticity field at the plane of symmetry that is perpendicular to the valve leaflets. The overall flow physics are very similar to those reported and discussed in great detail in [1] and for that they will not be discussed in detail herein. As shown in Fig. 16, during the opening phase, the flow field around the valve leaflets is organized and symmetric and is dominated by the vortex shedding from the leaflets and the valve housing. Past peak systole and throughout the closing stage, on the other hand, the flow in the vicinity to the valve is dominated by a number of small-scale chaotic vortices, which mark the transition to turbulence [1]. Overall the present FSI simulations yield flow fields that are in excellent agreement with the experimental measurements of [1] as shown in the Fig. 16 especially during the accelerating period of the systolic phase when the flow remains organized, laminar and repeatable from cycle to cycle and thus comparisons with instantaneous measurements are meaningful. In fact the overall level
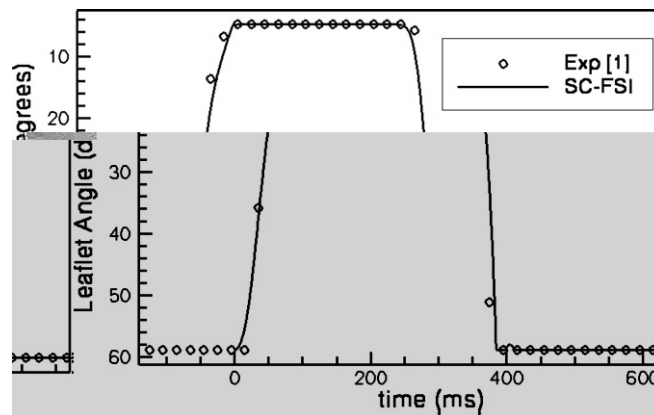

Fig. 15. BMHV simulation (fine mesh). Comparison of the calculated leaflet kinematics (solid line) with experimental observations [1] (circles).
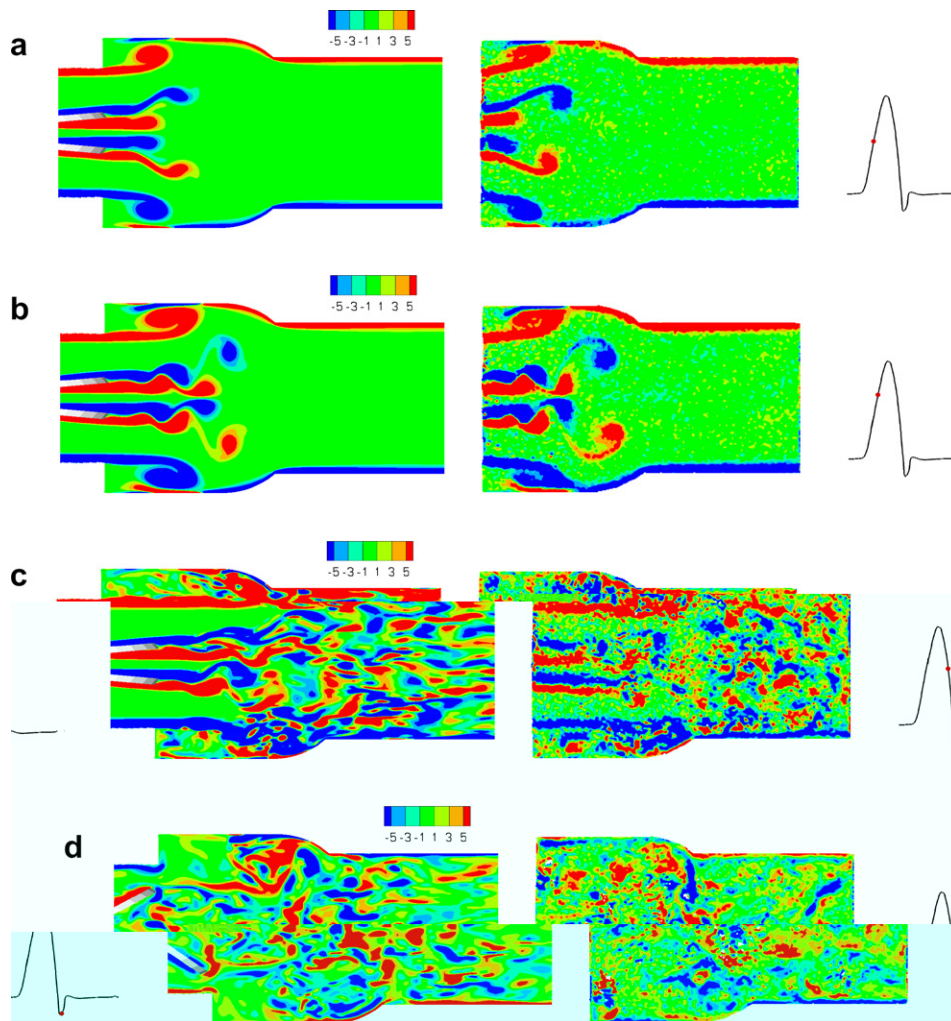
Fig. 16. BMHV simulation on fine grid (left) compared with the PIV measurements (right) of Dasi et al [1]. Instantaneous out-of-plane vorticity contours on the mid-plane of the valve at four different time instants within a cardiac cycle. The dots on the inflow waveform shown on the right of each figure indicate the time instant during the cycle. The contour levels are identical.

of agreement between measurements and simulations achieved by the present FSI simulation is similar and at many instances improved significantly compared to that achieved by Dasi et al. [1] who employed the same numerical method and grid resolution for the same geometry but with the leaflet kinematics prescribed from the experimental measurements. To further illustrate the enormous complexity of this FSI problem we visualize in Fig. 17 the instantaneous coherent structures shed from the valve leaflets and the aorta wall at several instants during the cardiac cycle. The vortical structures are visualized using the q-criterion [54] – see [1] for a detailed discussion of the underlying vorticity dynamics.

To explore the sensitivity of the computed results to mesh refinement we compare in Fig. 18 instantaneous vorticity iso-surfaces during the opening phase of the valve at three instants during the cardiac cycle. It is evident from this figure that even the coarse mesh captures with reasonable accuracy the coherent vorticity dynamics obtained on the fine mesh. As shown in Fig. 19, however, some appreciable discrepancies between the coarse and fine mesh solutions are observed in the predicted leaflet kinematics during the early opening stage. More specifically, on the coarse mesh the valve opening is initially delayed by 20 ms but at subsequent times the valve accelerates so that it ends up opening fully somewhat faster than in the fine mesh simulation. This discrepancy should be attributed to the fact that in the very early stages of valve opening, the FSI
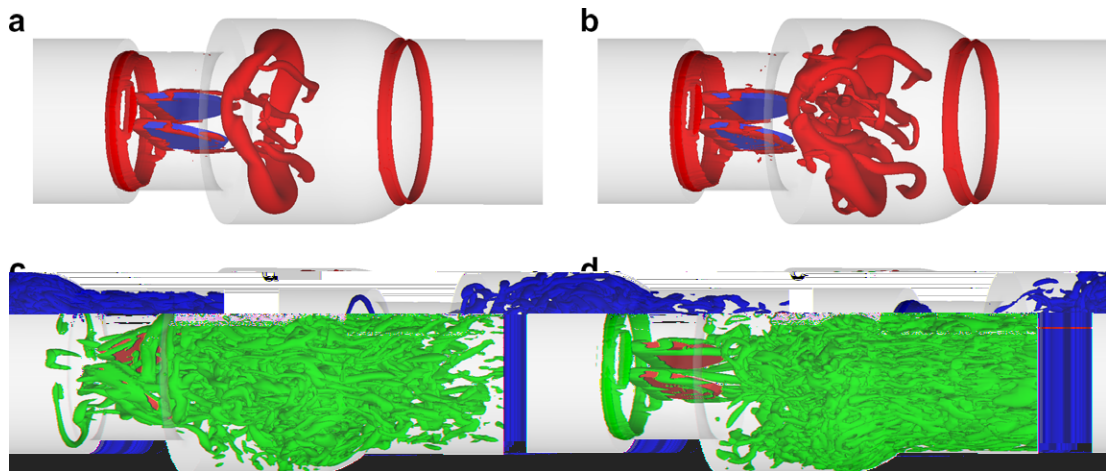
Fig. 17. BMHV simulation (fine grid). Instantaneous vortical structures visualized by iso-surfaces of q-criteria at the same four time instants as of Fig. 16 within a cardiac cycle.

dynamics is governed by the jets that emanate in the gap regions between the valve leaflets and the housing wall. As the valve leaflets begin to move forward from their closed position these gaps are initially small and a sufficiently fine mesh is required to resolve the flows through them accurately enough to yield the correct motion of the leaflet during the initial opening phase. Clearly the coarse mesh resolution is not fine enough to capture the early vorticity dynamics of the gap jets, which explains the delayed opening of the valve. These discrepancies not withstanding, however, the coarse mesh solution is in reasonable overall agreement with the solution obtained on the fine mesh.

Although the leaflets are geometrically symmetric and exposed to a symmetric incoming flow, the calculated leaflet kinematics is asymmetric as shown in Figs. 19 and 20. The asymmetry is very small and not detectable to plotting scale accuracy during the opening phase (Fig. 19) but it is very clear during the closing phase. This finding is consistent with and can be explained by the complexity of the flow field during the closing phase when, as shown in Fig. 16, the flow has broken down into a chaotic, turbulent-like state (see [1] for more details). These chaotic vortices lead to different moments on the two leaflets thus yielding the asymmetric motion. Asymmetric leaflet motion was also observed in the experiments [1]. Due to the nature of experiments, however, disturbances that can lead to asymmetric kinematics can arise from different sources, such as imperfection in the geometry of each leaflet, apparatus-specific noise and asymmetric disturbances at the inflow, etc. Such uncertainties make it essentially impossible to conclusively determine from experiments alone whether the asymmetric leaflet motion is caused by flow instabilities or other sources. In our numerical simulation, however, both the incoming conditions and the leaflet geometry are perfectly symmetric yet the computed leaflet motion is asymmetric. It can be argued that such asymmetry could be due to asymmetries inherent in the numerical algorithm. For example, possible sources of numerical asymmetries could be the QUICK upwind differencing used to discretize the convective terms and the incomplete convergence of iterative methods for solving the Poisson and the momentum equations. The impact of convergence-related asymmetries on the numerical solution is controllable by the convergence tolerance. In all our simulations the Poisson equation was iterated to machine zero, so incomplete convergence is unlikely to be the source of the observed asymmetries. In addition, comparing the coarse and fine mesh solutions (Fig. 18) suggests that assymetries due to the spatial discretization scheme are also unlikely to be the source of the asymmetric features of the computed solution. Therefore, it is reasonable to conclude that the observed asymmetries in the computed leaflet kinematics are real and are related to asymmetries in the flow due to natural flow instabilities. To the best of our knowledge this experimentally documented feature of BMHV flows has never been captured before computationally.

All previous results have been obtained with the SC-FSI algorithm, which yields efficient and robust solutions throughout the cardiac cycle. A striking finding of our simulations, however, is that even though during
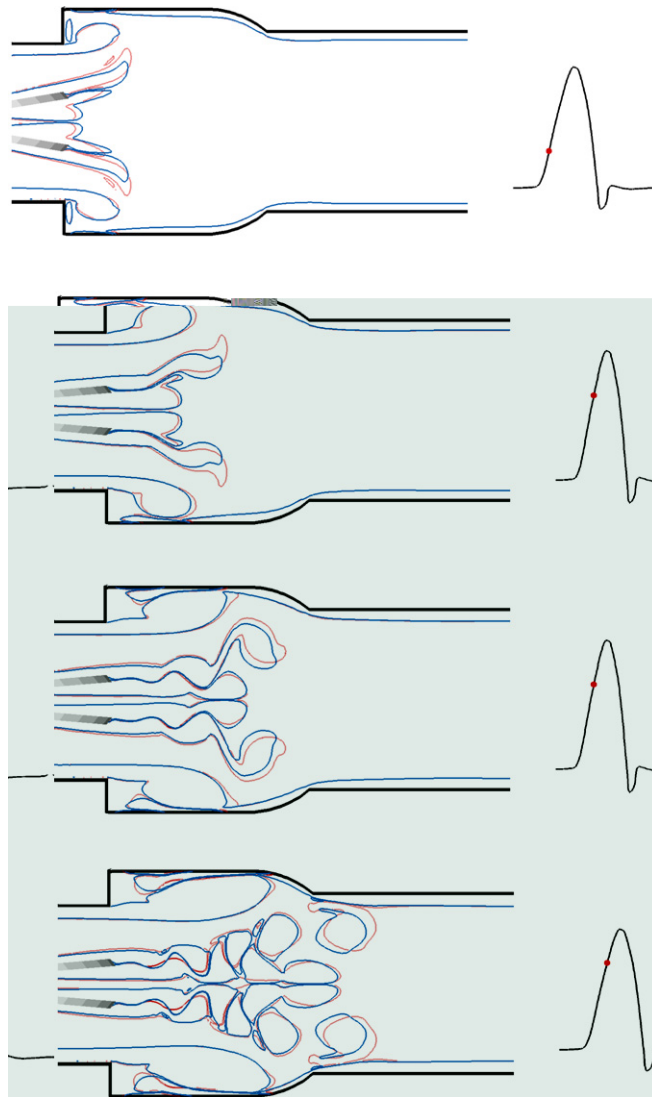
Fig. 18. BMHV simulation results on fine grid (solid blue lines) compared with the coarse grid (dotted red lines). Instantaneous out-of-plane vorticity iso-lines for vorticity magnitude equal to 1 on the mid-plane of the valve. The dots on the inflow waveform shown on the bottom of each figure column indicate the time instant during the cardiac cycle. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

the opening phase of the valve only the SC-FSI algorithm works, after the valve has opened fully and throughout the closing phase both the LC-FSI and SC-FSI are stable. In fact, as shown in Fig. 20 the two algorithms yield essentially identical results. A closer look into the two solutions is given in Fig. 21, which shows at a magnified scale the variations of the position θ, angular velocity ω, and moment coefficient $C_{\mathfrak{Z}}$ for one of the leaflets during the closing phase. When zooming to this scale, it is evident that the LC-FSI solution for θ exhibits small oscillations about the SC-FSI solution just before the leaflet comes to a complete stop. As expected the oscillations are very pronounced in the variation of the leaflet angular velocity and the moment. The reasons for this distinct oscillatory behavior will become clear in the subsequent section. The oscillations not withstanding, however, it is evident that the LC-FSI solution remains bounded throughout the closing phase and the computed leaflet position is very close to that obtained by SC-FSI. This finding is especially surprising when we note that the two algorithms work and yield similar results during the phase of the cardiac cycle when the flow is very complex and chaotic and the leaflets experience significantly higher acceleration.
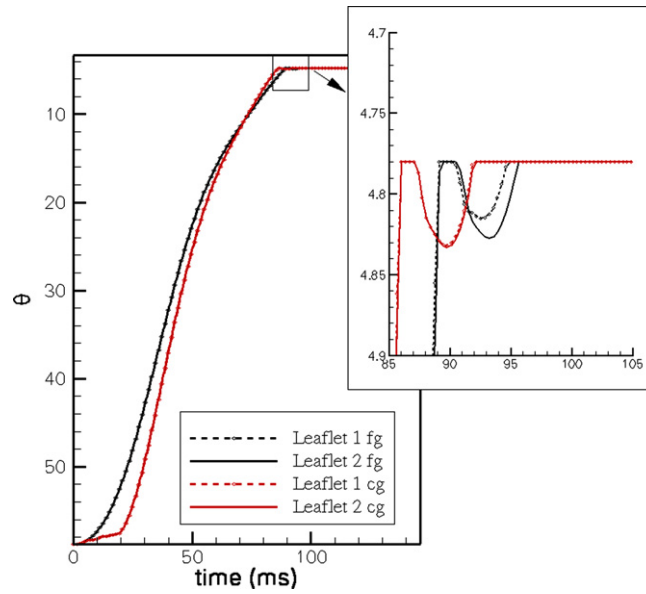
Fig. 19. BMHV simulation. Comparison of the upper and lower leaflet kinematics during the opening stage obtained by SC-FSI on the fine (black) and coarse (red) grids. fg and cg refer to fine and coarse grid, respectively. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)
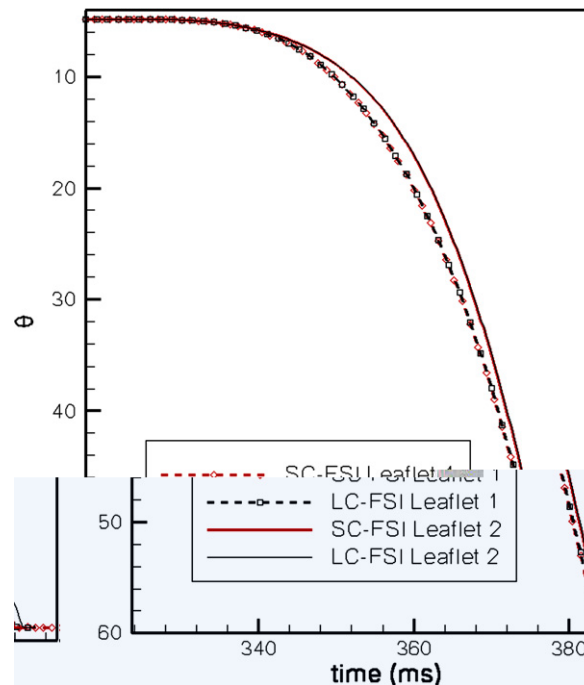

Fig. 20. BMHV simulation. Comparison of the upper and lower leaflet kinematics during the closing stage obtained by SC-FSI and LC-FSI.

Clearly in addition to the relative magnitude of the mass or inertia of the structure other parameters should be important in determining the stability of the LC-FSI algorithm.

Another striking finding is that the under-relaxation coefficient for the SC-FSI calculated by the Aitken's acceleration technique is generally larger in the closing phase than the opening phase. In fact and as seen in
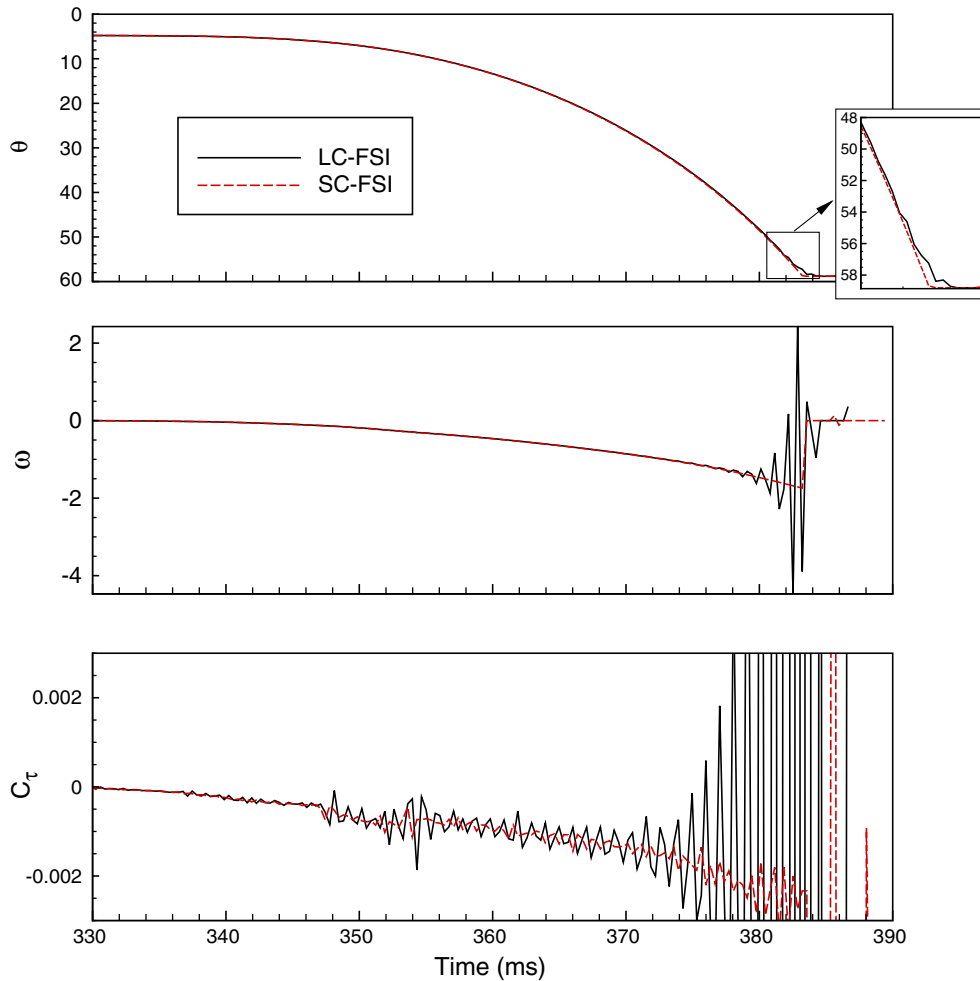
Fig. 21. BMHV simulation. Comparison of the LC-FSI vs. SC-FSI for one of the leaflets in the closing phase in terms of leaflets angle (top), angular velocity (middle), and moment coefficient (bottom).

Fig. 22 at many instances during closing the Aitken's acceleration technique over-relaxes the system yielding values of α greater than one. Once again and for the reasons discussed before, this is a rather surprising finding as at first glance the relative complexity of the flow would suggest that the closing phase should be computationally more challenging than the opening phase. These seemingly odd findings are explained in the subsequent Section 3.2.3 where we present a theoretical analysis of the stability of the various FSI algorithms.

3.2.3. Stability and convergence of FSI approaches

In this section we attempt to explain the stability and convergence properties of the LC-FSI and SC-FSI algorithms observed in the BMHV simulations by analyzing the stability of simpler, albeit similar, FSI problems. Rigorous analysis of the stability of the LC-FSI or SC-FSI algorithms requires the coupled consideration of the stability for both the fluid and structure equations. The structure can affect the stability of the fluid equations, Eq. (2), through the boundary condition in Eq. (4) while the fluid can affect the stability of the structure through the fluid imparted force term in the right-hand-side of the Eq. (3). Carrying out a formal stability analysis for the coupled FSI equations, however, is quite complicated and beyond the scope of this work. In this section we simplify the problem by considering only the stability of the structural equation with the effects of the fluid modeled through the added mass effect. Therefore, the stability of the discrete dynamic equation of the structure is analyzed not exactly in the form that it is solved in the FSI algorithm, Eq. (7) but
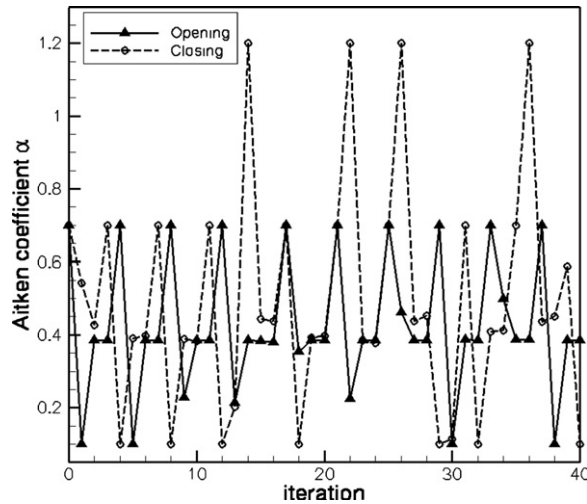
Fig. 22. Comparison of the under-relaxation coefficient α calculated by the Aitken's acceleration technique for SC-FSI during the early opening and closing phases.

in the form in which the added mass term has been decomposed from the rest of the force–see Eq. (13). For each FSI coupling approach, the added mass term in Eq. (13) is integrated in time in the same manner as the force term in Eq. (7), i.e., implicitly for SC-FSI and explicitly for LC-FSI. To begin our discussion, let us first summarize the key findings concerning FSI stability and convergence as observed in the results presented in the previous section:

(i) Oscillatory behavior of the computed solution: We found that for SC-FSI with constant under-relaxation coefficient the sub-iterations approach the final solution in a non-monotonic manner via a series of damped oscillations about the final solution. That is, if during one time step the computed solution overshoots the exact solution (positive error) it will under-shoot (negative error) the exact solution during the next time step and vice versa (see Fig. 14). Similarly during the end of the closing phase, the LC-FSI solution oscillates about the "exact" SC-FSI solution (see Fig. 20).

(ii) The angular velocity oscillations are much larger than the angle oscillations (see Fig. 20).

(iii) SC-FSI requires under-relaxation for convergence and stability, and the convergence rate depends on the under-relaxation coefficient α.

(iv) The LC-FSI iteration is stable during the closing phase but becomes rapidly unstable during opening.

(v) The SC-FSI under-relaxation coefficient obtained from the Aitken's acceleration technique was generally larger in closing phase in comparison with the opening phase.

To explain the above observations, let us first consider a simple case of a rigid plate immersed in stagnant ideal fluid. The plate starts rotating around a hinge by the action of a constant externally applied moment (see Fig. 23a). The plate's governing equation of motion can be written similar to the Eq. (13) with the added mass explicitly decomposed from the rest of the force:

$$\rho_s I \frac{\partial^2 \theta}{\partial t^2} = -\rho_f H \frac{\partial^2 \theta}{\partial t^2} + \mathfrak{I}^{\text{ext}} \tag{33}$$

where, $H > 0$ is the added mass constant, which depends on geometry, and $\mathfrak{I}^{\text{ext}}$ is the constant external moment around the hinge axis. Note that the term $\mathfrak{I}_{\text{flow}}$ is zero in this case since the fluid is assumed to be ideal and initially at rest. To solve Eq. (33) iteratively similar to the BMHV case, we discretize in the following manner:

$$\rho_s I \frac{\theta^{\ell+1} - 2\theta^n + \theta^{n-1}}{\Delta t^2} = -\rho_f H \frac{\theta^\ell - 2\theta^n + \theta^{n-1}}{\Delta t^2} + \mathfrak{I}^{\text{ext}} \tag{34}$$
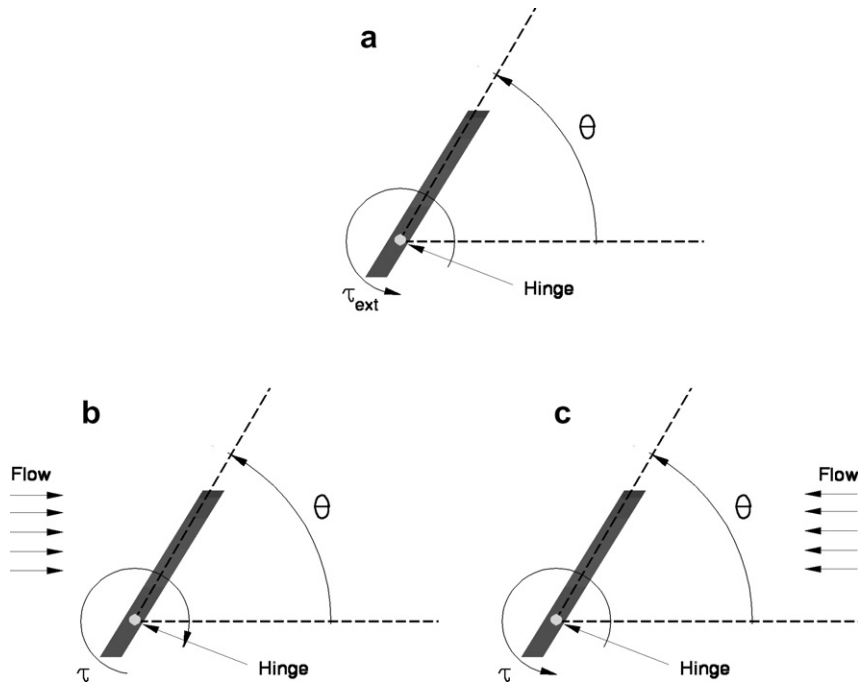
Fig. 23. A plate rotating around a hinge in still fluid (a) and (b), (c) in a free stream. The plate is rotated by (a) a constant external moment and (b), (c) the moment exerted by the flow. The configuration and flow direction in (b) is similar to the opening phase and (c) is similar to the closing phase.

For SC-FSI, the above equation is iterated over ℓ ($\ell = 0, L_{\max}$) until $\|\theta^{\ell+1} - \theta^\ell\| < \varepsilon$ and the solution converges to the solution of the discretized equation θ^{n+1}. For LC-FSI only one iteration is performed: $\theta^{\ell=0}$ is the initial approximation of θ^{n+1} based on an interpolation method and $\theta^{\ell=1}$ is the solution of the LC-FSI at $n + 1$ time level.

Proposition 1. *Added mass is responsible for the oscillatory behavior of the error.*

Proof. Assume that the θ^{n+1} is the exact solution of the discretized Eq. (34) at time $n + 1$. By introducing the error terms in Eq. (34), we obtain the following error equation:

$$\rho_s I \epsilon^{\ell+1} = -\rho_f H \epsilon^\ell \tag{35}$$

Therefore,

$$\epsilon^{\ell+1} = -\frac{\rho_f H}{\rho_s I} \epsilon^\ell \tag{36}$$

Noting that $\frac{\rho_f H}{\rho_s I}$ is always positive, the proposition is concluded that

$$\text{sign}(\epsilon^{\ell+1}) = -\text{sign}(\epsilon^\ell) \quad \square$$

Proposition 2. *For stability the ratio $\frac{\rho_f H}{\rho_s I}$ (called the FSI stability ratio hereafter) should be less than or equal to 1.*

Proof. For stability $|\frac{\epsilon^{\ell+1}}{\epsilon^\ell}| \leqslant 1$. Consequently, using Eq. (36) the proposition is reached that

$$\frac{\rho_f H}{\rho_s I} \leqslant 1 \quad \square \tag{37}$$

Remark 3. In the absence of flow and in an ideal fluid the stability of the FSI coupling is independent of the time step, i.e., if the stability ratio $\frac{\rho_f H}{\rho_s I} > 1$ then both SC-FSI (without under-relaxation) and LC-FSI are unstable, no matter how small the time step.

Remark 4. The FSI stability condition Eq. (37) is similar to the stability condition found in [33,35] for an internal FSI problem of flow inside a domain with compliant walls.

Remark 5. The stability of the scheme not only depends on the fluid to structure density ratio but also the geometry of the structure.

Remark 6. Reducing the reduced mass M_{red} in the VIV cylinder case reduced the structure to fluid density ratio ρ_s/ρ_f, which, consequently, violated the FSI stability condition Eq. (37).

Remark 7. In the BMHV case, the stability condition Eq. (37) is violated due to the small inertia I (geometry).

Proposition 8. *The angular velocity error is $O(1/\Delta t)$ larger than the angle error.*

Proof. Using the second order accurate upwind scheme to calculate the angular velocity at $\ell + 1$, we have:

$$\omega^{\ell+1} = \frac{3\theta^{\ell+1} - 4\theta^n + \theta^{n-1}}{2\Delta t} \tag{38}$$

Similarly, the exact solution for angular velocity ω at $n + 1$ will be:

$$\omega^{n+1} = \frac{3\theta^{n+1} - 4\theta^n + \theta^{n-1}}{2\Delta t} \tag{39}$$

Subtracting Eq. (38) from Eq. (39), the error equation is obtained:

$$\epsilon_\omega^{\ell+1} = \frac{3\epsilon^{\ell+1}}{2\Delta t} \tag{40}$$

Therefore,

$$\frac{\epsilon_\omega^{\ell+1}}{\epsilon^{\ell+1}} = O(1/\Delta t) \tag{41}$$

and the proposition is reached. $\square$

Proposition 9. *The angular velocity solution exhibits an oscillatory behavior similar to the angle solution.*

Proof. Similar to the Eq. (40), the error for at level ℓ will be:

$$\epsilon_\omega^\ell = \frac{3\epsilon^\ell}{2\Delta t} \tag{42}$$

Therefore, dividing Eq. (40) by Eq. (42) and using Eq. (36), we get:

$$\frac{\epsilon_\omega^{\ell+1}}{\epsilon_\omega^\ell} = \frac{\epsilon^{\ell+1}}{\epsilon^\ell} = -\frac{\rho_f H}{\rho_s I} \tag{43}$$

Consequently,

$$\text{sign}(\epsilon_\omega^{\ell+1}) = -\text{sign}(\epsilon_\omega^\ell)$$

and the proposition is reached. $\square$

Remark 10. Considering the small time step in BMHV simulations along with Propositions 8 and 9 the observation (ii) listed at the beginning of this section is clarified.

Proposition 11. *If $\frac{\rho_f H}{\rho_s I} > 1$ then under-relaxation is necessary for the convergence of the strong coupling with*

$$\frac{2}{1 + \frac{\rho_f H}{\rho_s I}} > \alpha > 0$$

Proof. The error equation for the under-relaxation Eq. (14) is:

$$\tilde{\epsilon}^{\ell+1} = (1-\alpha)\tilde{\epsilon}^{\ell} + \alpha\epsilon^{\ell+1} \tag{44}$$

and similar to Eq. (36),

$$\epsilon^{\ell+1} = -\frac{\rho_f H}{\rho_s I}\tilde{\epsilon}^{\ell} \tag{45}$$

Therefore,

$$\tilde{\epsilon}^{\ell+1} = \left[-\alpha\frac{\rho_f H}{\rho_s I} + (1-\alpha)\right]\tilde{\epsilon}^{\ell} \tag{46}$$

For convergence $\left|\frac{\tilde{\epsilon}^{\ell+1}}{\tilde{\epsilon}^{\ell}}\right| < 1$, and consequently, $\left|-\alpha\frac{\rho_f H}{\rho_s I} + (1-\alpha)\right| < 1$. After some mathematical manipulations the proposition is obtained. □

Remark 12. This explains why when the under-relaxation coefficient is close to one the SC-FSI iteration fails to converge.

By the above simple case, i.e., rigid plate rotation in still fluid, the observations (i),(ii), and (iii) have been explained. However, to explain the observations (iv) and (v) the flow effects should be considered as well.

The simplest case containing flow is a rigid plate free to rotate around a fixed hinge axis in a free stream (see Fig. 23b,c). The moment $\mathfrak{I}_{\text{flow}}$ on the plate in this case is a function of the angle of the plate relative to the free stream θ. When the plate is parallel to the steam the moment $\mathfrak{I}_{\text{flow}}$ will be zero and when the plate is perpendicular to the flow the moment is maximum. To examine the stability of FSI in this case, a similar approach as in the previous analysis can been adopted. Nevertheless, the moment here is not constant and changes with θ, i.e., $\mathfrak{I}_{\text{flow}} = \mathfrak{I}_{\text{flow}}(\theta)$.

Proposition 13. *For stability*

$$\left|-\frac{\rho_f H}{\rho_s I} + \frac{\Delta t^2}{\rho_s I}\frac{\partial \mathfrak{I}_{\text{flow}}}{\partial \theta}\right|$$

should be less than or equal to 1.

Proof. The error equation can be obtained in a similar manner as proposition 1 but recognizing that $\mathfrak{I}_{\text{flow}}^{n+1} = \mathfrak{I}_{\text{flow}}^{\ell} + \frac{\partial \mathfrak{I}_{\text{flow}}}{\partial \theta}\epsilon^{\ell}$:

$$\rho_s I\epsilon^{\ell+1} = -\rho_f H\epsilon^{\ell} + \Delta t^2\frac{\partial \mathfrak{I}_{\text{flow}}}{\partial \theta}\epsilon^{\ell} \tag{47}$$

$$\epsilon^{\ell+1} = \left(-\frac{\rho_f H}{\rho_s I} + \frac{\Delta t^2}{\rho_s I}\frac{\partial \mathfrak{I}_{\text{flow}}}{\partial \theta}\right)\epsilon^{\ell} \tag{48}$$

For stability $\left|\frac{\epsilon^{\ell+1}}{\epsilon^{\ell}}\right| \leqslant 1$. Therefore,

$$\left|-\frac{\rho_f H}{\rho_s I} + \frac{\Delta t^2}{\rho_s I}\frac{\partial \mathfrak{I}_{\text{flow}}}{\partial \theta}\right| \leqslant 1 \tag{49}$$

and the proposition is reached. □

Remark 14. Depending on the sign of the term $\frac{\partial \mathfrak{I}_{\text{flow}}}{\partial \theta}$ the stability constraint Eq. (49) can be either more or less restrictive than the stability constraint Eq. (37) in the no flow case. Noting that the FSI stability ratio $\frac{\rho_f H}{\rho_s I}$ is always positive, then if $\frac{\partial \mathfrak{I}_{\text{flow}}}{\partial \theta}$ is negative it adds to the added mass effect. Therefore, the constraint (49) will be more restrictive and the flow becomes unstable at a lower FSI stability ratio. Decreasing the time step here has a stabilizing effect since it reduces the flow effect. However, if the FSI stability ratio $\frac{\rho_f H}{\rho_s I} \geqslant 1$, decreasing the time step has no effect since the ratio (49) is already larger than unity. Similarly, if the $\frac{\partial \mathfrak{I}_{\text{flow}}}{\partial \theta}$ is positive it reduces the added mass effect. Therefore, the constraint (49) will be less restrictive and the flow becomes unstable at a

higher FSI stability ratio. Surprisingly, decreasing the time step has a destabilizing effect here since it reduces the stabilizing effect of the flow.

Remark 15. Looking at the configuration Fig. 23b, which is similar to the opening phase, it is observed that that the moment decreases as the plate moves, i.e., $\frac{\partial \Im_{\text{flow}}}{\partial \theta} < 0$. Therefore, the flow adds to the added mass effect and has a destabilizing effect.

Remark 16. Looking at the configuration Fig. 23c, which is similar to the closing phase, it is observed that the moment increases as the plate moves, i.e., $\frac{\partial \Im_{\text{flow}}}{\partial \theta} > 0$. Therefore, the flow reduces the added mass effect and has a stabilizing effect.

Remark 17. Based on Remarks 15 and 16, the opening is less stable than closing. Therefore, the observations (iv) and (v) has been clarified.

4. Summary and conclusions

We extended the CURVIB approach to simulate FSI problems involving multiple, arbitrarily complex solid bodies undergoing arbitrarily large deformations. Major algorithmic advances reported in this work are: (1) the development of an efficient ray-tracing algorithm for quickly identifying the location of the moving immersed bodies in the background grid; and (2) the development of a strong coupling FSI algorithm that is stable, robust and efficient even for systems with very low mass for which loose-coupling iterations fail to yield converged solutions. The profound effect of under-relaxation and in particular of the Aitken's acceleration technique on the stability and efficiency of the SC-FSI algorithm was also demonstrated conclusively. We applied the CURVIB-FSI algorithm to simulate the VIV of elastically mounted 2D cylinders and physiologic pulsatile flow through a clinical quality BMHV. For both cases the computed results were shown to be in excellent agreement with benchmark numerical results and experimental measurements. We also presented a theoretical stability analysis for generic FSI problems to explore the role of the added mass term and the local flow properties on the stability and convergence of the FSI iterations. With this simple analysis we were able to explain all findings from the numerical experiments and derive the required upper bound for the under-relaxation coefficient of stable SC-FSI iterations.

A major conclusion from this work is that the stability and convergence of the SC-FSI and LC-FSI algorithms depends both on the properties of the structure (mass and geometry) and the sign of the local temporal variation of the force or moment imparted by the flow on the structure. In general, for FSI problems involving structural components with very low inertia (i.e., when the added mass effects dominate the FSI dynamics) LC-FSI is very difficult to converge and it is often unstable regardless of the size of the time step. The same can be true for the SC-FSI algorithm. The stability problems of the SC-FSI algorithm, however, can be effectively eliminated by using under-relaxation and in particular the Aitken's acceleration technique, which determines the optimal relaxation factor in a dynamic manner as part of the solution. SC-FSI with Aitken's acceleration is thus shown to be a powerful technique for simulating complex FSI problems even for system with very low inertia.

A striking conclusion from the BMHV simulation, which is characterized by the very low inertia of the valve leaflets, is that the LC-FSI algorithm is, as one would expect, unstable during the valve opening phase but the same algorithm is stable and yields very similar solution to that obtained by the SC-FSI algorithm during the valve closing phase. This is a particularly surprising finding when one takes into account the relative complexity of the instantaneous flow during the opening and closing phases – the flow being well organized and laminar during the opening phase and very complex, turbulent-like during valve closing. Our stability analysis clarified this seemingly paradoxical finding by showing that during the closing phase the sign of the rate of change of the flow-induced moment on the valve leaflets is such that flow effects tend to alleviate the adverse stability effects of the added mass term, which arise due to the very low inertia of the valve leaflets. These results point to the conclusion that the decision of whether LC-FSI or SC-FSI is used for a particular problem should take into account both the structural and local flow properties. For the BMHV case, for

example, the entire closing phase can be simulated with the LC-FSI algorithm, which greatly enhances the overall efficiency of the simulation (recall that unlike SC-FSI, LC-FSI does not require costly FSI sub-iterations to obtain a converged solution every time step).

To the best of our knowledge the BMHV work reported in this paper is the first ever to tackle and successfully solve the coupled FSI problem in a full 3D manner at physiologic Reynolds numbers and using high numerical resolution and realistic valve leaflet properties. Our future work will focus on extending the CUR-VIB-FSI algorithm to FSI problems involving flexible surfaces so that the more complex problem of the native trileaflet aortic heart valve can be tackled at the same resolution as the present BMHV simulations.

Acknowledgments

References

[1] L.P. Dasi, L. Ge, H.A. Simon, F. Sotiropoulos, A.P. Yoganathan, Vorticity dynamics of a bileaflet mechanical valve in an axisymmetric aorta, Physics of Fluids 19 (2007) 067105.
[2] L. Ge, F. Sotiropoulos, A numerical method for solving the 3D unsteady incompressible Navier–Stokes equations in curvilinear domains with complex immersed boundaries, Journal of Computational Physics 225 (2007) 1782–1809.
[3] J. Donea, S. Giuliani, J. Halleux, An arbitrary Lagrangian–Eulerian finite element method for transient dynamic fluid–structure interactions, Computer Methods in Applied Mechanics and Engineering 33 (1–3) (1982) 689–723.
[4] C.S. Peskin, Flow patterns around heart valves: a numerical method, Journal of Computational Physics 10 (1972) 252–271.
[5] C. Farhat, M. Lesoinne, P. LeTallec, Load and motion transfer algorithms for fluid/structure interaction problems with non-matching discrete interfaces: Momentum and energy conservation, optimal discretization and application to aeroelasticity, Computer Methods in Applied Mechanics and Engineering 157 (1) (1998) 95–114.
[6] P. Le Tallec, J. Mouro, Fluid structure interaction with large structural displacements., Computer Methods in Applied Mechanics and Engineering 190 (24) (2001) 3039–3067.
[7] M. Fernandez, M. Moubachir, A Newton method using exact jacobians for solving fluid–structure coupling, Computers and Structures 83 (2–3) (2005) 127–142.
[8] R. Dillon, L. Fauci, A. Fogelson, D. Gaver III, Modeling biofilm processes using the immersed boundary method, Journal of Computational Physics 129 (1) (1996) 57–73.
[9] J. Stockie, S. Green, Simulating the motion of flexible pulp fibres using the immersed boundary method, Journal of Computational Physics 147 (1998) 147–165.
[10] L. Zhu, C. Peskin, Simulation of a flapping flexible filament in a flowing soap film by the immersed boundary method, Journal of Computational Physics 179 (2) (2002) 452–468.
[11] E. Fadlun, R. Verzicco, P. Orlandi, J. Mohd-Yusof, Combined immersed-boundary finite-difference methods for three-dimensional complex flow simulations, Journal of Computational Physics 161 (1) (2000) 35–60.
[12] J. Kim, D. Kim, H. Choi, An immersed-bounday finite-volume method for simulations of flow in complex geometries, Journal of Computational Physics 171 (2001) 132–150.
[13] Y. Tseng, J. Ferziger, A ghost-cell immersed boundary method for flow in complex geometry, Journal of Computational Physics 192 (2) (2003) 593–623.
[14] S. Marella, S. Krishnan, H. Liu, H. Udaykumar, Sharp interface Cartesian grid method I: an easily implemented technique for 3D moving boundary computations, Journal of Computational Physics 210 (1) (2005) 1–31.
[15] J. Choi, R. Oberoi, J. Edwards, J. Rosati, An immersed boundary method for complex incompressible flows, Journal of Computational Physics 224 (2) (2007) 757–784.
[16] D. Le, B. Khoo, J. Peraire, An immersed interface method for viscous incompressible flows involving rigid and flexible boundaries, Journal of Computational Physics 220 (1) (2006) 109–138.
[17] S. Xu, Z. Wang, An immersed interface method for simulating the interaction of a fluid with moving boundaries, Journal of Computational Physics 216 (2) (2006) 454–493.
[18] A. Gilmanov, F. Sotiropoulos, E. Balaras, A general reconstruction algorithm for simulating flows with complex 3D immersed boundaries on Cartesian grids, Journal of Computational Physics 191 (2) (2003) 660–669.
[19] A. Gilmanov, F. Sotiropoulos, A hybrid cartesian/immersed boundary method for simulating flows with 3D, geometrically complex, moving bodies, Journal of Computational Physics 207 (2005) 457–492.
[20] H. Udaykumar, R. Mittal, W. Shyy, Computation of solid–liquid phase fronts in the sharp interface limit on fixed grids, Journal of Computational Physics 153 (2) (1999) 535–574.

[21] R. Leveque, Z. Li, The immersed interface method for elliptic equations with discontinuous coefficients and singular sources, SIAM Journal on Numerical Analysis 31 (4) (1994) 1019–1044.

[22] R. Mittal, G. Iaccarino, Immersed boundary methods, Annual Review of Fluid Mechanics 37 (2005) 239–261.

[23] D. Kim, H. Choi, Immersed boundary method for flow around an arbitrarily moving body, Journal of Computational Physics 212 (2) (2006) 662–680.

[24] K. Schulz, Y. Kallinderis, Unsteady flow structure interaction for incompressible flows using deformable hybrid grids, Journal of Computational Physics 143 (2) (1998) 569–597.

[25] C. Farhat, M. Lesoinne, Two efficient staggered algorithms for the serial and parallel solution of three-dimensional nonlinear transient aeroelas-tic problems, Computer Methods in Applied Mechanics and Engineering 182 (3) (2000) 499–515.

[26] C. Farhat, K. van Der Zee, P. Geuzaine, Provably second-order time-accurate loosely-coupled solution algorithms for transient nonlinear computational aeroelasticity, Computer Methods in Applied Mechanics and Engineering 195 (17–18) (2006) 1973–2001.

[27] H. Matthies, J. Steindorf, Partitioned strong coupling algorithms for fluid–structure interaction, Computers and Structures 81 (8–11) (2003) 805–812.

[28] H. Ahn, Y. Kallinderis, Strongly coupled flow/structure interactions with a geometrically conservative ALE scheme on general hybrid meshes, Journal of Computational Physics 219 (2) (2006) 671–696.

[29] G. Tryggvason, B. Bunner, A. Esmaeeli, N. Al-Rawahi, W. Tauber, J. Han, Y. Jan, D. Juric, S. Nas, A front-tracking method for the computations of multiphase flow, Journal of Computational Physics 169 (2) (2001) 708–759.

[30] J. Sethian, Level Set Methods and Fast Marching Methods, Cambridge University Press, Cambridge, 1999.

[31] C. Hirt, B. Nichols, Volume of fluid (VOF) method for the dynamics of free boundaries, Journal of Computational Physics 39 (1) (1981) 201–225.

[32] E. Haines, Point in polygon strategies, Academic Press Graphics Gems Series, 1994 (pp. 24–46)..

[33] P. Causin, J. Gerbeau, F. Nobile, Added-mass effect in the design of partitioned algorithms for fluid-structure problems, Computer Methods in Applied Mechanics and Engineering 194 (42–44) (2005) 4506–4527.

[34] C. Conca, A. Osses, J. Planchard, Added mass and damping in fluid–structure interaction, Computer Methods in Applied Mechanics and Engineering 146 (3) (1997) 387–405.

[35] C. Forster, W. Wall, E. Ramm, Artificial added mass instabilities in sequential staggered coupling of nonlinear structures and incompressible viscous flows, Computer Methods in Applied Mechanics and Engineering 196 (7) (2007) 1278–1293.

[36] D.

[38] E. Spanier, Algebraic Topology, Springer-Verlag, 1981.

[39] T. Moller, B. Trumbore, Fast, minimum storage ray-triangle intersection, Journal of Graphics Tools 2 (1) (1997) 21–28.

[40] C. Williamson, R. Govardhan, Vortex-induced vibrations, Annual Review of Fluid Mechanics 36 (2004) 413–455.

[41] R.D. Blevins, Flow-Induced Vibration, second ed., Van Nostrand Rein-hold, New York, NY, 1990.

[42] T. Sarpkaya, A critical review of the intrinsic nature of vortex-induced vibrations, Journal of Fluids and Structures 19 (4) (2004) 389.

[43] C.H.K. Williamson, R. Govardhan, Vortex-induced vibrations, Annual Review of Fluid Mechanics 36 (2004) 413.

[44] G.Moe, Moe, Z.J.Wu, Wu, Lift force on a vibrating cylinder in a current, Proceedings of the International Offshore Mechanics and Arctic Engineering Symposium, vol. 2, American Society of Mechanical Engineers (ASME), New York, NY, USA, Hague, Netherlands, 1989, p. 259 (compilation and indexing terms, Copyright 2006 Elsevier Inc. All rights reserved 89120412524 Lift Forces Vibrating Forces).

[45] L.L. Pauley, P. Moin, W.C. Reynolds, The structure of two-dimensional separation, Journal of Fluid Mechanics 220 (1990) 397–411.

[46] A. Yoganathan, K. Chandran, F. Sotiropoulos, Flow in prosthetic heart valves: state-of-the-art and future directions, Annals of Biomedical Engineering 33 (12) (2005) 1689–1694.

[47] A.P. Yoganathan, Z. He, S.C. Jones, Fluid mechanics of heart valves, Annual Review of Biomedical Engineering 6 (2004) 331–362.

[48] Y. Lai, K. Chandran, J. Lemmon, A numerical simulation of mechanical heart valve closure fluid dynamics, Journal of Biomechanics 35 (7) (2002) 881–892.

[49] J. Vierendeels, K. Dumont, E. Dick, P. Verdonck, Analysis and stabilization of fluid-structure interaction algorithm for rigid-body motion, AIAA Journal 43 (12) (2005) 2549–2557.

[50] S. Krishnan, H. Udaykumar, J. Marshall, K. Chandran, Two-dimensional dynamic simulation of platelet activation during mechanical heart valve closure, Annals of Biomedical Engineering 34 (10) (2006) 1519–1534.

[51] G. Pedrizzetti, F. Domenichini, Flow-driven opening of a valvular leaflet, Journal of Fluid Mechanics 569 (2006) 321–330.

[52] R. Cheng, Y. Lai, K. Chandran, Three-dimensional fluid-structure interaction simulation of bileaflet mechanical heart valve flow dynamics, Annals of Biomedical Engineering 32 (11) (2004) 1471–1483.

[53] J. De Hart, G. Peters, P. Schreurs, F. Baaijens, A three-dimensional computational analysis of fluid–structure interaction in the aortic valve, Journal of Biomechanics 36 (1) (2003) 103–112.

[54] J.C.R. Hunt, A.A. Wray, P. Moin, Eddies, streams, and convergence zones in turbulent flows, in: Proceedings of the 1988 Summer Program on Studying Turbulence Using Numerical Simulation Databases (SEE N89-24538 18–34), vol. 2, 1988, pp. 193–208.