

Document Version

Final published version

Licence

CC BY

Citation (APA)

Li, J., Ji, Y., Verhelst, H., den Besten, H., & Möller, M. (2027). Parameterization-driven arbitrary Lagrangian–Eulerian method for large-deformation isogeometric fluid–structure interaction. *Computer Methods in Applied Mechanics and Engineering*, 463, Article 119358. <https://doi.org/10.1016/j.cma.2026.119358>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

In case the licence states “Dutch Copyright Act (Article 25fa)”, this publication was made available Green Open Access via the TU Delft Institutional Repository pursuant to Dutch Copyright Act (Article 25fa, the Taverne amendment). This provision does not affect copyright ownership.
Unless copyright is transferred by contract or statute, it remains with the copyright holder.

Sharing and reuse

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.



Parameterization-driven arbitrary Lagrangian–Eulerian method for large-deformation isogeometric fluid–structure interaction

Jingya Li ^a, Ye Ji ^{a,*}, Hugo Verhelst ^b, Henk den Besten ^c, Matthias Möller ^a

^a Delft Institute of Applied Mathematics, Delft University of Technology, Delft, 2628 CD, The Netherlands

^b Department of Mechanical Engineering, Eindhoven University of Technology, Eindhoven, 5612 AE, The Netherlands

^c Maritime and Transport Technology Department, Delft University of Technology, Delft, 2628 CD, The Netherlands

ARTICLE INFO

Keywords:

Isogeometric analysis
Fluid–structure interaction
Arbitrary Lagrangian–Eulerian (ALE)
Multi-patch spline parameterization
Large-deformation mesh motion

ABSTRACT

Body-fitted arbitrary Lagrangian–Eulerian (ALE) methods provide a sharp representation of the fluid–structure interface but rely on mesh-update strategies that incrementally deform a reference configuration. Because each update is applied on top of the previous one, distortion accumulates over time, and large structural motion eventually leads to loss of mesh validity. This path dependence is inherent to the mesh-motion formulation itself: it follows from updating each mesh from the previous one, and cannot be eliminated by refining the underlying extension operator. To address this issue, we reformulate the ALE mesh-motion problem in the isogeometric setting as a sequence of independent domain parameterization problems. At each time step, a multi-patch spline parameterization of the fluid domain is constructed from the current interface geometry; its validity and quality therefore depend on the current configuration alone, not on the cumulative history of mesh deformation. Mesh motion follows directly from this per-step parameterization, rather than being the primary unknown of an auxiliary mesh-extension problem driven by interface displacement, so the cumulative distortion characteristic of classical ALE is absent by construction. Three technical components realize this framework: (i) a barrier-function-based spline parameterization that enforces a strictly positive Jacobian at every time step; (ii) a tangential-slip reparameterization that handles sustained large rotations of closed interfaces, where no fixed boundary-to-parameter correspondence is admissible; and (iii) a constant-preserving quasi-interpolation operator for solution transfer between consecutive parameterizations, which preserves uniform states in the adopted advective ALE update. The framework is formulated in a partitioned setting, so that the fluid and structural solvers remain independent. Our numerical examples start with a purely kinematic example motivating the use of higher-order spline-based mesh formulations over conventional piecewise-linear methods in terms of the maximum allowed deformation the mesh can handle. Thereafter, we validate our method in combination with higher-order isogeometric mesh descriptions on two two-dimensional FSI benchmarks and a prescribed-rotation flow benchmark, covering standard and large-rotation regimes, and on three-dimensional falling-sphere and rotor problems. On the rotating-square benchmark, the tangential-slip strategy sustains continuous rotation beyond the admissible range of conventional mesh-deformation ALE methods with a fixed boundary correspondence. On the Turek–Hron and perpendicular-flap benchmarks, the method reproduces published reference displacements and responses, while the minimum scaled Jacobian remains bounded away from zero throughout the simulation. A prescribed-translation

* Corresponding author.

E-mail addresses: j.li-9@tudelft.nl (J. Li), y.ji-1@tudelft.nl (Y. Ji), h.m.verhelst@tue.nl (H. Verhelst), Henk.denBesten@tudelft.nl (H.d. Besten), m.moller@tudelft.nl (M. Möller).

<https://doi.org/10.1016/j.cma.2026.119358>

Received 30 April 2026; Received in revised form 28 August 2026; Accepted 29 August 2026

Available online 9 September 2026

0045-7825/© 2026 The Authors. Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

sphere test isolates mesh robustness under a 40 m displacement, and a fully coupled free-fall test reproduces the wall-corrected terminal velocity to within 0.29%. A three-dimensional rotor example further demonstrates that the framework extends naturally to volumetric spline parameterizations. Finally, we show that the per-step spline parameterizations can be used directly within a standard finite element solver, yielding results consistent with the isogeometric solution on the same sequence of geometries. This decoupling of geometry construction from field discretization positions the proposed framework as a solver-agnostic geometry component, compatible with both isogeometric and classical finite-element FSI pipelines. The focus is therefore on the geometry-construction module, rather than on claiming general accuracy or efficiency advantages of isogeometric analysis over finite-element remeshing strategies.

1. Introduction

Many problems in computational fluid dynamics are posed on domains whose geometry evolves in time. These include free-surface and multiphase flows [1], flows past moving or rotating bodies, and fluid–structure interaction (FSI) problems [2,3]. In all of these settings, the computational mesh must be updated at every time step to conform to the moving boundary, and the validity of this update determines the range of deformations that the simulation can sustain. FSI is particularly demanding in this respect: the fluid domain is not prescribed a priori but follows from the coupled fluid–structure response, so the mesh must track an interface that can accumulate large displacement or rotation over the course of the simulation.

FSI methods can be broadly classified by how they represent the fluid–structure interface [4,5]. *Non-body-fitted* approaches discretize the fluid equations on a background grid that does not conform to the interface, avoiding the difficulties of a deforming mesh. This class includes immersed boundary methods [6], CutFEM and related unfitted formulations [7,8], and their isogeometric counterparts [9,10]. Their flexibility, however, comes at the price of a non-conforming interface representation, which typically requires special techniques to recover near-wall quantities such as interface tractions, the very quantities that drive the FSI coupling. *Body-fitted* approaches align the computational mesh with the interface and provide a sharp geometric representation of the boundary; for FSI, this sharpness is not an optional refinement but a direct enabler of accurate interface loads [11,12].

Within the body-fitted class, the arbitrary Lagrangian–Eulerian (ALE) formulation [13,14] is the most widely used and also the oldest framework. In ALE, mesh nodes on the interface move with the structure, while interior nodes are repositioned by an auxiliary mesh-motion strategy, most commonly a harmonic [15], biharmonic [16], or linear-elastic extension of the interface displacement [17,18]. Because the mesh itself evolves in time, any ALE discretization must additionally comply with the discrete geometric conservation law, a sufficient condition for the nonlinear stability of the time integration [19,20].

This framework has since been transferred to the isogeometric setting [21], giving rise to a family of IGA-ALE formulations in which spline and NURBS bases replace piecewise-polynomial shape functions for both geometry representation and solution approximation [22–25]. Beyond the familiar benefits of geometric exactness at curved interfaces and higher-order inter-element continuity, IGA-ALE offers one feature that is especially pertinent to FSI coupling: the interface tractions emerge as naturally smooth fields, obviating the post-processing usually required to reconstruct them from piecewise-polynomial stresses [11,12]. The present work is set within this IGA-ALE framework and adopts a partitioned coupling strategy: the fluid and structural subproblems are solved by independent solvers that exchange only interface displacements and tractions at each coupling step.

The principal limitation of classical body-fitted ALE methods lies in the mesh-motion problem itself. In standard mesh-deformation formulations, the mesh at time t^{n+1} is obtained from the mesh at time t^n by applying an incremental displacement field, typically the solution of an auxiliary extension problem driven by the interface displacement. The resulting update is path-dependent: the mesh at any instant is the composition of all previous updates, and a local distortion introduced at one step is carried forward to every subsequent step. For moderate interface motion this accumulation remains bounded, but under large cumulative displacement or rotation it is not: the Jacobian determinant of the ALE mapping eventually passes through zero, and the mesh becomes invalid. The top row of Fig. 1 illustrates this behavior for a rotating-square configuration.

Three concrete consequences follow. First, element distortion degrades the interface tractions that drive the coupled problem. Second, the Jacobian determinant of the ALE mapping is not guaranteed to remain positive: once it changes sign at any point of the mesh, the mapping is no longer a bijection and the simulation has to be stopped. Third, on distorted meshes the discrete geometric conservation law can no longer be enforced, and the associated nonlinear stability guarantees [19] are lost. Remeshing and mesh-smoothing strategies can mitigate these effects [26], but the underlying cause persists: as long as the mesh at t^{n+1} is constructed from the mesh at t^n , any refinement of the extension operator stays within the same path-dependent construction.

Several strategies have been proposed to push body-fitted ALE beyond the regime of small-to-moderate interface motion. Fixed-mesh ALE formulations [27,28] avoid mesh motion altogether by solving the governing equations on a fixed background mesh, retaining the ALE mapping only in a virtual sense; this sidesteps the distortion problem but sacrifices the sharp body-fitted interface that motivates the ALE formulation itself. A more direct line of work keeps the body-fitted character of ALE but replaces mesh deformation with per-step parameterization: rather than updating the existing mesh, a new conforming mesh is generated at every time step from the current interface configuration. Recent work by Schwenner and Fries [29] and by Schalk and Fries [30] shows that this strategy is viable for large-deformation FSI in the finite element setting, and Di Cristofaro et al. [31] demonstrate a related ALE approach with frequent remeshing and field remapping for three-dimensional FSI with large boundary motion, including




Fig. 1. Two approaches to ALE mesh motion under large interface displacement, illustrated on a rotating-square configuration. **Top:** in the classical mesh-deformation ALE approach, the mapping $\mathcal{A}_{t^{n+1}}$ is obtained by composing the preceding mapping $\mathcal{A}_{t^n}$ with an incremental update $\Delta \mathcal{A}^{n+1}$; the resulting path-dependent distortion may accumulate until the mesh becomes invalid. **Bottom:** in the proposed parameterization-driven approach, the spline geometry $\mathcal{G}^{n+1} := \{\mathbf{F}_k^{n+1}\}_{k=1}^{n_p}$ is reconstructed directly from the current interface configuration $\partial\Omega^f(t^{n+1})$, independently of $\mathcal{G}^n$. The small triangular marker on the inner square is used to track its cumulative rotation.

flapping-wing and falling-sphere configurations. Within this line of work, the transfer of discrete fields between successive meshes emerges as a central concern; constant-preserving transfer operators [32,33] are of particular interest because they preserve uniform states across the transfer. For conservative ALE formulations, this transfer property must be complemented by a discrete geometric identity for evolving Jacobian and mesh flux [19].

In the isogeometric setting, ALE-FSI formulations have so far been developed primarily within the classical mesh-deformation approach [22,25]. While these formulations benefit from the geometric exactness and higher-order continuity of spline discretizations, they inherit the structural limitation identified in the previous section: the fluid mesh at any time instant is still the product of a cumulative sequence of incremental updates.

An alternative approach becomes available when the fluid discretization is based on spline representations of the geometry. The present work is aligned with this mesh-regeneration viewpoint, but its emphasis is different from conventional finite-element remeshing. The regenerated object is a spline parameterization of the fluid domain. This makes it possible to control validity directly through the positivity of the parameterization Jacobian, to introduce tangential slip at the level of boundary pre-images, and to reuse the same geometry update in both isogeometric and finite element flow solvers.

In isogeometric analysis, the physical domain is represented not by a collection of nodal coordinates, but by a spline- or NURBS-based parameterization that maps a fixed parametric domain onto the physical domain. This distinction is more than representational. A finite element mesh is tied to its construction history: its quality depends on how it was obtained, and subsequent updates can only reposition its existing nodes. A spline parameterization, in contrast, is a self-contained geometric object that can be rebuilt from scratch at any time instant, using the current interface configuration as its sole input. The bottom row of Fig. 1 illustrates this property on the same rotating-square configuration: at every snapshot, the mesh is an independently constructed parameterization rather than the accumulated result of all previous updates.

We therefore reformulate the ALE mesh-motion problem. Instead of updating the fluid mesh incrementally from one time step to the next, we construct it at each step as an independent spline parameterization of the current fluid domain, using the current interface configuration as input. The past motion of the interface is recorded by the structural solver alone; the fluid mesh at time t^n does not depend on the one at t^{n-1} . The mesh velocity that couples the fluid equations to the moving domain is then recovered from successive parameterizations rather than prescribed through an auxiliary mesh-extension problem.

This reformulation has consequences that reach beyond robustness under large deformations. The cumulative distortion typical of classical ALE does not arise at all since each parameterization is built independently. The solution transfer between consecutive parameterizations is performed by a constant-preserving quasi-interpolation operator, which preserves uniform states in the adopted advective ALE update. While this transfer property alone does not establish the fully discrete geometric conservation law, the latter is recovered when it is combined with a consistent discrete relation between the evolving Jacobian and the mesh flux. Consequently, the discrete geometric conservation law holds algebraically, and the associated nonlinear stability guarantees [19] are retained. Finally, because the parameterization depends only on the interface geometry and not on the fluid discretization, the same sequence of parameterizations can drive either an isogeometric or a standard finite element fluid solver. The purpose of the present work is therefore not to claim a universal accuracy or efficiency advantage of IGA over finite-element remeshing strategies.

Instead, the contribution is a spline-parameterization-driven ALE geometry module whose validity is controlled at the level of the parameterization and which can be reused by different downstream flow discretizations. To the best of our knowledge, this view of ALE mesh motion has not been pursued in the isogeometric FSI literature, even though the required spline parameterization machinery for complex domains is now available [34].

Our main contributions are as follows.

- We reformulate the ALE mesh-motion problem in the isogeometric setting as a sequence of independent domain parameterization problems. At every time step, we build a new multi-patch spline parameterization of the fluid domain directly from the current interface, rather than updating the previous mesh. This replaces the path-dependent mesh-deformation approach, and therefore the cumulative distortion it carries, with an effectively history-independent construction.
- For closed domains that rotate through large cumulative angles, a fixed boundary-to-parameter correspondence simply cannot be maintained. Once the body has rotated far enough, it becomes incompatible with any valid mesh. We address this by letting the parametric pre-images slide tangentially along the interface. The resulting scheme handles rotation regimes that, to our knowledge, cannot be reached by any mesh-deformation ALE formulation.
- We use a constant-preserving quasi-interpolation operator to transfer the discrete solution between successive parameterizations. In the adopted advective ALE update, exact reproduction of constants yields free-stream preservation.
- Because each parameterization is built as a pure geometric object, nothing in the construction is tied to the field solver. We exploit this by feeding the same sequence of parameterizations into a standard finite element solver and recovering results consistent with the isogeometric computation.

The remainder of the paper is organized as follows. Section 2 states the governing equations of the fluid and structural subproblems together with the corresponding ALE formulation. Section 3 develops the numerical framework: the isogeometric discretization of the fluid equations, the per-step spline parameterization of the fluid domain together with the tangential-slip reparameterization for closed configurations, the recovery of the ALE mesh velocity from two consecutive parameterizations, and the constant-preserving solution-transfer operator used for free-stream preservation in the adopted advective ALE update. Section 4 reports two- and three-dimensional numerical experiments: standard FSI benchmarks, a rotating-square configuration that drives the inner body through a full revolution, prescribed-translation and freely falling spheres, a three-dimensional rotor problem, and a test in which the regenerated parameterizations are coupled to a classical finite element solver. Section 6 summarizes the main findings and outlines directions for future work.

2. Governing equations

In this paper, we consider a fluid–structure interaction problem in which the fluid occupies a time-dependent domain $\Omega^f(t) \subset \mathbb{R}^d$ ($d \in \{2, 3\}$), separated from the surrounding structure $\Omega^s(t)$ by a moving interface $\Gamma_{fs}(t) = \partial\Omega^f(t) \cap \partial\Omega^s(t)$. Throughout, superscripts f and s distinguish fluid and structural quantities, and the subscript 0 denotes quantities in the reference configuration.

2.1. Incompressible flow

The fluid motion is governed by the incompressible Navier–Stokes equations in the time-dependent domain $\Omega^f(t)$:

$$\rho^f \left(\frac{\partial \mathbf{u}^f}{\partial t} + (\mathbf{u}^f \cdot \nabla) \mathbf{u}^f \right) = \nabla \cdot \boldsymbol{\sigma}^f + \mathbf{f}^f \quad \text{in } \Omega^f(t), \quad (1a)$$

$$\nabla \cdot \mathbf{u}^f = 0 \quad \text{in } \Omega^f(t), \quad (1b)$$

where $\mathbf{u}^f : \Omega^f(t) \times [0, T] \rightarrow \mathbb{R}^d$ is the fluid velocity, ρ^f is the (constant) density, $\mathbf{f}^f$ is a body force, and $\boldsymbol{\sigma}^f$ is the Cauchy stress tensor.

For an incompressible Newtonian fluid, the stress tensor is given by

$$\boldsymbol{\sigma}^f = -p^f \mathbf{I} + 2\mu^f \mathbb{D}(\mathbf{u}^f), \quad (2)$$

with pressure p^f , dynamic viscosity $\mu^f = \rho^f \nu^f$, and symmetric strain-rate tensor $\mathbb{D}(\mathbf{u}^f) = \frac{1}{2}(\nabla \mathbf{u}^f + \nabla \mathbf{u}^{f\top})$.

The system is closed by an initial condition $\mathbf{u}^f(\cdot, 0) = \mathbf{u}_0^f$ in $\Omega^f(0)$ together with Dirichlet and Neumann conditions on the inflow and outflow boundaries; the interface conditions on $\Gamma_{fs}(t)$ are deferred to Section 2.4.

2.2. Arbitrary Lagrangian–Eulerian formulation

The time-dependence of $\Omega^f(t)$ requires a formulation of the moving-grid. To this end, we adopt the arbitrary Lagrangian–Eulerian (ALE) framework. A reference configuration Ω_0^f is mapped onto the current fluid domain at each time instant through the application of the following:

$$\mathcal{A} : \Omega_0^f \times [0, T] \rightarrow \mathbb{R}^d, \quad \mathbf{x} = \mathcal{A}(\mathbf{X}, t) = \mathbf{X} + \mathbf{u}^{\text{ALE}}(\mathbf{X}, t), \quad (3)$$




Fig. 2. Configurations and mappings in the ALE-FSI framework, shown on a representative conforming parameterization in which every node on the interface belongs to both the fluid and the structural parameterization (open circles). **Left:** reference configurations of the fluid (Ω_0^f) and the structure (Ω_0^s), sharing the reference interface $\Gamma_{fs,0}$; $\mathbf{X}$ denotes a reference interface point, and $\mathbf{n}_0^s$ is the outward unit normal on the solid. **Right:** current configurations at time t , obtained from the reference by the ALE mapping $\mathcal{A}_t$ on the fluid side and by the structural deformation $\mathbf{X} \mapsto \mathbf{X} + \mathbf{d}^s(\mathbf{X}, t)$ on the solid side. The two mappings agree on $\Gamma_{fs,0}$, so the current interface $\Gamma_{fs}(t)$ is the same curve from the fluid and structural sides. The point $\mathbf{x} = \mathcal{A}_t(\mathbf{X})$ lies on $\Gamma_{fs}(t)$ with outward normals $\mathbf{n}^s$ and $\mathbf{n}^f = -\mathbf{n}^s$.

where $\mathbf{u}^{\text{ALE}}$ is the mesh displacement field, and $\Omega^f(t) := \mathcal{A}(\Omega_0^f, t)$. The mesh velocity is the time derivative taken at fixed reference point $\mathbf{X}$,

$$\hat{\mathbf{w}}(\mathbf{X}, t) := \left. \frac{\partial \mathcal{A}}{\partial t} \right|_{\mathbf{X}} = \left. \frac{\partial \mathbf{u}^{\text{ALE}}}{\partial t} \right|_{\mathbf{X}}, \quad (4)$$

which pushes forward to the current configuration as $\mathbf{w}(\mathbf{x}, t) = \hat{\mathbf{w}}(\mathcal{A}^{-1}(\mathbf{x}, t), t)$.

Under the ALE mapping, the time derivative in the momentum equation is evaluated at fixed $\mathbf{X}$ rather than fixed $\mathbf{x}$. The fluid equations become

$$\rho^f \left(\left. \frac{\partial \mathbf{u}^f}{\partial t} \right|_{\mathbf{X}} + ((\mathbf{u}^f - \mathbf{w}) \cdot \nabla) \mathbf{u}^f \right) = \nabla \cdot \boldsymbol{\sigma}^f + \mathbf{f}^f \quad \text{in } \Omega^f(t), \quad (5a)$$

$$\nabla \cdot \mathbf{u}^f = 0 \quad \text{in } \Omega^f(t). \quad (5b)$$

The quantity $\mathbf{u}^f - \mathbf{w}$ is the relative velocity of the fluid with respect to the moving mesh.

On the interface, consistency of the ALE mapping with the structural motion requires

$$\mathbf{u}^{\text{ALE}}(\mathbf{X}, t) = \mathbf{d}^s(\mathbf{X}, t) \quad \text{on } \Gamma_{fs,0} \times [0, T], \quad (6)$$

where $\mathbf{d}^s$ is the structural displacement defined in Section 2.3. Within the interior of the fluid domain, $\mathbf{u}^{\text{ALE}}$ is not uniquely determined and must be constructed such that the ALE mapping is locally invertible. Denoting the Jacobian matrix of $\mathcal{A}$ by

$$\mathbf{J}^{\text{ALE}}(\mathbf{X}, t) := \nabla_{\mathbf{X}} \mathcal{A}(\mathbf{X}, t), \quad (7)$$

this reads

$$\det \mathbf{J}^{\text{ALE}}(\mathbf{X}, t) > 0 \quad \text{for all } (\mathbf{X}, t) \in \Omega_0^f \times [0, T]. \quad (8)$$

The discretization strategy used to ensure (8) is developed in Section 3.2.1; the corresponding discrete criterion is stated in Section 3.5.

Fig. 2 summarizes the reference and current configurations together with the mappings that connect them on the fluid and structural sides.

2.3. Structural mechanics

The structural response is described in a total Lagrangian framework with respect to a reference configuration $\Omega_0^s \subset \mathbb{R}^d$. The unknowns are the displacement field $\mathbf{d}^s : \Omega_0^s \times [0, T] \rightarrow \mathbb{R}^d$ and the corresponding velocity $\mathbf{v}^s := \partial \mathbf{d}^s / \partial t$. Balance of linear momentum reads

$$\rho_0^s \frac{\partial^2 \mathbf{d}^s}{\partial t^2} = \nabla_0 \cdot \mathbf{P}^s + \rho_0^s \mathbf{b}^s \quad \text{in } \Omega_0^s \times [0, T], \quad (9)$$

where ρ_0^s is the reference mass density, $\mathbf{b}^s$ is a body force per unit mass, ∇_0 denotes the gradient with respect to reference coordinates, and $\mathbf{P}^s := \mathbf{F} \mathbf{S}^s$ is the first Piola–Kirchhoff stress tensor.

The deformation gradient is defined by $\mathbf{F} := \mathbf{I} + \nabla_0 \mathbf{d}^s$, and the Green–Lagrange strain tensor by

$$\mathbf{E} := \frac{1}{2} (\mathbf{F}^T \mathbf{F} - \mathbf{I}), \quad (10)$$

which captures geometric nonlinearity under large deformations. The structure is modeled as a St. Venant–Kirchhoff material, for which the second Piola–Kirchhoff stress is

$$\mathbf{S}^s = \lambda^s \operatorname{tr}(\mathbf{E}) \mathbf{I} + 2\mu^s \mathbf{E}, \quad (11)$$

where λ^s, μ^s are the Lamé parameters. The associated Cauchy stress, used in the interface traction condition below, is obtained by the standard push-forward,

$$\boldsymbol{\sigma}^s = J^{-1} \mathbf{F} \mathbf{S}^s \mathbf{F}^T, \quad J := \det \mathbf{F}. \quad (12)$$

The reference boundary $\partial\Omega_0^s$ is partitioned into three non-overlapping parts: a Dirichlet part Γ_D^s , a Neumann part Γ_N^s , and the fluid–structure interface $\Gamma_{fs,0} \subset \partial\Omega_0^s$. On $\Gamma_{fs,0}$, the traction is prescribed by the fluid and is treated in Section 2.4. On the remaining parts,

$$\mathbf{d}^s = \mathbf{g}^s \quad \text{on } \Gamma_D^s \times [0, T], \quad (13a)$$

$$\mathbf{P}^s \mathbf{n}_0^s = \mathbf{h}^s \quad \text{on } \Gamma_N^s \times [0, T], \quad (13b)$$

where $\mathbf{g}^s$ is a prescribed displacement, $\mathbf{h}^s$ is a prescribed reference traction, and $\mathbf{n}_0^s$ is the outward unit normal on the reference configuration. Initial data $\mathbf{d}^s(\cdot, 0) = \mathbf{d}_0^s$ and $\mathbf{v}^s(\cdot, 0) = \mathbf{v}_0^s$ close the structural subproblem.

2.4. Coupling conditions

At each time t , the fluid and structural subproblems are linked through conditions on the interface $\Gamma_{fs}(t)$. The partitioned formulation solves each subproblem independently and uses these conditions to transfer information at each coupling step [12].

Kinematic coupling. Continuity of velocity at the interface requires

$$\mathbf{u}^f(\mathbf{x}, t) = \mathbf{v}^s(\mathbf{X}, t) \quad \text{for } \mathbf{x} = \mathbf{X} + \mathbf{d}^s(\mathbf{X}, t), \quad \mathbf{X} \in \Gamma_{fs,0}, \quad (14)$$

where $\mathbf{v}^s = \partial \mathbf{d}^s / \partial t$ is the structural velocity. Consistently, the ALE mapping tracks the structural displacement on the interface:

$$\mathbf{u}^{\text{ALE}}(\mathbf{X}, t) = \mathbf{d}^s(\mathbf{X}, t) \quad \text{on } \Gamma_{fs,0} \times [0, T]. \quad (15)$$

Condition (15) determines the ALE mesh displacement on $\Gamma_{fs,0}$; its extension into the interior is the subject of Section 3.2.1.

Dynamic coupling. Equilibrium of tractions across the interface reads

$$\boldsymbol{\sigma}^f(\mathbf{x}, t) \mathbf{n}^f + \boldsymbol{\sigma}^s(\mathbf{x}, t) \mathbf{n}^s = \mathbf{0} \quad \text{on } \Gamma_{fs}(t), \quad (16)$$

where $\mathbf{n}^f$ and $\mathbf{n}^s = -\mathbf{n}^f$ are the outward unit normals on the current fluid and structural interfaces. Because the structural subproblem is posed in the reference configuration, (16) is recast using Nanson's formula $J \mathbf{F}^{-T} \mathbf{n}_0^s dA_0 = \mathbf{n}^s dA$, which yields the reference-configuration form

$$\mathbf{P}^s \mathbf{n}_0^s = -J \boldsymbol{\sigma}^f \mathbf{F}^{-T} \mathbf{n}_0^s \quad \text{on } \Gamma_{fs,0} \times [0, T]. \quad (17)$$

Eq. (17) acts as a Neumann boundary condition for the structural subproblem: the traction applied by the fluid is evaluated on the deformed interface and pulled back to the reference configuration.

3. Parameterization-driven ALE framework

This section develops the numerical framework that realizes the parameterization-driven ALE formulation. Four technical components make up the core of the framework: the isogeometric discretization of the fluid equations (Section 3.1); the per-step spline parameterization of the fluid domain (Section 3.2); the ALE mesh velocity recovered from successive parameterizations (Section 3.3); and the constant-preserving solution transfer that secures the discrete geometric conservation law (Section 3.4). A discrete criterion for the bijectivity of the resulting patch mappings is collected in Section 3.5.

3.1. Isogeometric discretization for fluid subproblem

We discretize the fluid subproblem using isogeometric analysis (IGA) [21], with the specific choices detailed below: a multi-patch NURBS representation of the geometry and a Taylor–Hood-type B-spline pair for velocity and pressure.

B-spline and NURBS basis. For a given open knot vector $\Xi = \{\xi_1 \leq \xi_2 \leq \dots \leq \xi_{n+p+1}\}$, the univariate B-spline basis $\{B_{i,p}\}_{i=1}^n$ of degree p is constructed by the Cox–de Boor recursion [21,35]. Tensor-product B-spline bases on the parametric domain $\mathcal{P} = [0, 1]^d$ are obtained as products of univariate bases in each parametric direction. Introducing positive weights yields the corresponding NURBS basis functions $\{N_i(\xi)\}_{i=1}^n$, and the geometry of each patch is represented as

$$\mathbf{x}(\xi) = \sum_{i=1}^n N_i(\xi) \mathbf{x}_i, \quad \xi \in \mathcal{P}, \quad (18)$$

where $\{\mathbf{x}_i\}_{i=1}^n \subset \mathbb{R}^d$ are the control points.

Multi-patch discretization. The fluid domain $\Omega^f(t)$ is represented as a collection of n_p conforming patches (see Section 3.2), with patch k parameterized by a time-dependent geometry map

$$\mathbf{F}_k(\xi, t) = \sum_{i=1}^{n_k} N_i(\xi) \mathbf{x}_{k,i}(t), \quad \mathbf{F}_k(\cdot, t) : \mathcal{P} \rightarrow \Omega_k(t), \quad (19)$$

where $\{\mathbf{x}_{k,i}(t)\}$ are the control points of patch k at time t . The discrete solution spaces are constructed by composition with $\mathbf{F}_k$, and global spaces are assembled by identifying control points shared by adjacent patches. At time t^n , the multi-patch spline geometry is denoted by

$$\mathcal{G}^n := \{\mathbf{F}_k^n\}_{k=1}^{n_p}. \quad (20)$$

Taylor–Hood discretization. To satisfy the inf-sup stability condition required for a stable discretization of the incompressible Navier–Stokes equations, we adopt a Taylor–Hood-type isogeometric velocity–pressure pair [36]. The velocity space V_h consists of B-splines of degree p and regularity C^{p-1} , while the pressure space Q_h consists of B-splines of degree $p-1$ and regularity C^{p-2} , both defined on the same physical element partition. This pair is inf-sup stable and delivers optimal-order approximation.

For $p \geq 2$, the C^{p-1} regularity of the velocity space makes the velocity gradient globally continuous on each patch. The viscous stress field σ^f is therefore continuous within each patch, and interface tractions can be evaluated directly at the patch boundary without the projection or recovery procedures required on C^0 Lagrange elements [11,12].

3.2. Per-step spline domain parameterization

Using the multi-patch setting of Section 3.1, the patch mappings $\{\mathbf{F}_k(\cdot, t)\}_{k=1}^{n_p}$ are constructed anew at each time step from the current boundary data, without reference to a deforming background mesh. This construction has two ingredients: a barrier-patch optimization procedure that produces a bijective, well-shaped parameterization from given boundary curves (Section 3.2.1); and two application patterns—one for the open fluid domains arising in FSI coupling (Section 3.2.2), and one for the closed annular domains used in our prescribed-motion benchmark (Section 3.2.3). In both settings the ingredient of Section 3.2.1 is applied patch-wise; the two settings differ only in how the patch boundaries are assembled from the interface and from fixed geometry.

3.2.1. Barrier-patch method

The barrier-patch method [34,37] builds a bijective and well-shaped NURBS parameterization of a single patch from prescribed boundary control points. Let $\mathbf{F}_k$ be the patch map of patch k , with boundary control points fixed by the surrounding geometry (specified separately for open and closed domains in Sections 3.2.2 and 3.2.3). The interior control points $\{\mathbf{x}_{k,i}\}_{i \in \mathcal{I}_k}$ are determined by a two-stage optimization, where $\mathcal{I}_k$ indexes the interior degrees of freedom of patch k .

Stage 1: Bijection at quadrature points. Let

$$\mathbf{J}_k(\xi) := \frac{\partial \mathbf{F}_k}{\partial \xi}(\xi), \quad \xi \in \mathcal{P}, \quad (21)$$

denote the Jacobian matrix of the patch map, and let $\{\xi_q\}_{q=1}^{n_q} \subset \mathcal{P}$ be a set of sampling points (taken as the Gauss quadrature points of the element-wise integration rule). We first eliminate fold-over configurations by minimizing

$$E_{\text{fold}}(\mathbf{F}_k) = \sum_{q=1}^{n_q} [\delta - \det \mathbf{J}_k(\xi_q)]_+^2, \quad (22)$$

where $[x]_+ := \max(x, 0)$ denotes the positive part and $\delta > 0$ is a small threshold. The functional (22) is strictly positive whenever the discrete bijectivity condition $\det \mathbf{J}_k(\xi_q) \geq \delta$ fails at any sampling point, and vanishes otherwise. Minimization is therefore halted once a configuration with $E_{\text{fold}} = 0$ is reached.

Stage 2: Quality improvement. Once bijectivity is secured, element quality is improved by minimizing the Winslow-plus-area functional

$$E_{\text{qual}}(\mathbf{F}_k) = \lambda_1 \int_{\mathcal{P}} \frac{\|\mathbf{J}_k\|_F^2}{\det \mathbf{J}_k} d\xi + \lambda_2 \int_{\mathcal{P}} (\det \mathbf{J}_k)^2 d\xi, \quad (23)$$

where $\|\cdot\|_F$ denotes the Frobenius norm. The first term is the classical Winslow functional, which penalizes angular distortion and acts as a smoother on the control net; the second term promotes uniform cell areas. The positivity of $\det \mathbf{J}_k$ established in Stage 1 is preserved throughout Stage 2 because the Winslow integrand becomes singular as $\det \mathbf{J}_k \rightarrow 0^+$. The weights $\lambda_1, \lambda_2 > 0$ are taken as $\lambda_1 = 1$ and $\lambda_2 = 1/A^2$, where $A = \int_{\mathcal{P}} \det \mathbf{J}_k d\xi$ is the (scaled) physical area. Section 4.

Both stages are minimized over the interior control points using L-BFGS, with the boundary control points held fixed. The integrals in (23) are evaluated by the same element-wise Gauss rule used in (22). The two-stage procedure returns a patch map $\mathbf{F}_k$ whose discrete bijectivity is certified by Stage 1 and whose element quality has been improved by Stage 2; the same procedure is applied independently to each patch at each time step.

3.2.2. Application to open FSI domains

For fluid domains arising in fluid–structure interaction, the fluid occupies the region exterior to a deforming body and interior to fixed walls. We partition the fluid domain into two layers:

- a *near-body layer* of n_s patches, each carrying one segment of the fluid–structure interface $\Gamma_{fs}(t)$ as its inner boundary;
- an *extension layer* of n_{ext} patches that connects the near-body layer to the fixed outer walls (inflow, outflow, and far-field boundaries).

Only the near-body patches are reparameterized at each time step; the extension patches are fixed once and for all at initialization. The two layers are joined conformally, and the control points on the shared interface between the layers are held fixed throughout the simulation.

Near-body patch boundaries. For near-body patch k , the boundary control points are assembled from three sources:

- the *inner* boundary $\gamma_k^{\text{in}}(\cdot, t^n)$ tracks the corresponding segment of $\Gamma_{fs}(t^n)$ and is supplied by the structural subproblem through the ALE consistency condition (15);
- the *outer* boundary γ_k^{out} lies on the fixed interface with the extension layer and is therefore time-independent;
- the *side* boundaries are shared with the two neighboring near-body patches and are constructed once at initialization, together with the rest of the near-body layer, so that only their inner boundaries evolve with the fluid–structure interface.

Given these four boundary curves, the interior control points of $\mathbf{F}_k$ are obtained by the barrier-patch procedure of Section 3.2.1.

Interface transfer. At each time step, the structural solver returns the updated positions of interface control points on $\Gamma_{fs}(t^n)$. In the spline-based coupling setting adopted in literature [11,12], the fluid and structural solvers exchange information through NURBS/spline representation of the interface. The inner boundary curves $\{\gamma_k^{\text{in}}(\cdot, t^n)\}$ are assembled from these control points and passed to the barrier-patch optimization as fixed boundary data.

3.2.3. Application to closed domains: ruled surface and tangential slip

For closed annular configurations, in which a simply connected body is fully surrounded by fluid, the fluid domain is discretized using $n_p = n_s$ conforming patches without extension layer. The inner and outer boundaries of the annulus are represented as closed B-spline curves $\gamma^{\text{in}}(\cdot, t)$ and $\gamma^{\text{out}}(\cdot)$, both parameterized over $s \in [0, 1]$ with periodic identification $s = 0 \sim s = 1$.

Ruled surface initialization. An initial global parameterization of the annular domain is constructed as the ruled surface between the two boundary curves,

$$\mathbf{F}^{\text{init}}(s, \eta, t) = (1 - \eta) \gamma^{\text{out}}(s) + \eta \gamma^{\text{in}}(s, t), \quad (s, \eta) \in [0, 1]^2, \quad (24)$$

which is then partitioned into n_s conforming patches by subdividing the parameter domain along s . The barrier-patch procedure of Section 3.2.1 is applied patch-wise to improve element quality; boundary control points on $\gamma^{\text{in}}(\cdot, t)$ and $\gamma^{\text{out}}(\cdot)$ remain fixed, while interior control points are updated.

Tangential slip. The ruled-surface construction (24) imposes a fixed parametric correspondence $s \mapsto s$ between inner and outer boundaries. Under prescribed rigid-body rotation of the inner body, this correspondence accumulates over time: a material point on the inner boundary rotates with the body, while the corresponding outer point remains stationary, producing patches that shear more and more with each step. After a finite number of steps, the patch maps become highly skewed and eventually non-bijective, and the barrier stage of Section 3.2.1 can no longer recover a valid parameterization [38].

We remove this accumulation by reparameterizing the inner boundary at each time step before constructing the ruled surface. Let $\delta^n \in [0, 1)$ denote a normalized tangential shift chosen (below) to align the inner boundary with the outer boundary. The shifted inner boundary is

$$\tilde{\gamma}^{\text{in}}(s, t^n) := \gamma^{\text{in}}((s + \delta^n) \bmod 1, t^n), \quad s \in [0, 1], \quad (25)$$

so that the physical curve is unchanged, while the location of the parametric seam is moved. In the implementation this is done by splitting the closed spline curve at the shifted parameter value and then merging the two pieces in the opposite order. The ruled surface (24) is then constructed with $\tilde{\gamma}^{\text{in}}$ instead of γ^{in} .

For the prescribed rigid-body rotations considered in this work, the shift is chosen directly from the cumulative rotation angle. Let θ^n be the rotation angle at time t^n . We use

$$\delta^n = \left(-\frac{\theta^n}{2\pi} \right) \bmod 1. \quad (26)$$

The sign is chosen so that the parameterization slips in the direction opposite to the physical rotation. In this way, the correspondence between the inner and outer boundaries remains close to the initial one instead of accumulating tangential shear.

For piecewise-smooth inner boundaries with n_c geometric corners distributed evenly in parameter, the corner parameters are

$$C = \left\{ \frac{j}{n_c} : j = 0, 1, \dots, n_c - 1 \right\} \quad (27)$$

($n_c = 4$ in the rotating-square and rotor examples of Sections 4.5 and 5). When δ^n falls close to an element of C , the split-and-merge step cuts $\gamma^{\text{in}}(\cdot, t^n)$ very near a genuine corner, producing a reparameterized curve with a spurious C^0 point next to the real corner. We

avoid this by displacing δ^n away from corners. Let d^n denote the signed periodic distance from δ^n to the nearest corner parameter. When $|d^n| < \varepsilon_c$, we apply the smooth local offset

$$\delta^n \leftarrow \delta^n + \varepsilon_s \left[1 - \left(\frac{d^n}{\varepsilon_c} \right)^2 \right]^2, \quad \varepsilon_c = 0.04, \quad \varepsilon_s = 0.025, \quad (28)$$

followed by reduction modulo 1. The values satisfy $\varepsilon_c < 1/(2n_c)$ and $\varepsilon_s < \varepsilon_c$, so that bump supports of adjacent corners do not overlap and the correction cannot push δ^n past the nearest corner. The bump $[1 - (d^n/\varepsilon_c)^2]^2$ is the lowest-order polynomial that vanishes to first order at $|d^n| = \varepsilon_c$, ensuring that δ^n depends continuously on θ^n .

As a final safeguard, when δ^n is within machine precision of 0 or 1, the split-and-merge step reduces to the identity up to round-off and is skipped. The reparameterization (25)–(28) enables the construction of well-shaped meshes under arbitrarily large cumulative tangential motion; this is demonstrated on the rotating-square benchmark in Section 4.5.

Scope of the tangential-slip strategy. The tangential-slip construction used here assumes that the closed interface admits a stable cyclic boundary correspondence. This is the case for convex or nearly convex bodies such as the rotating square and the rotor geometry considered later. For strongly non-convex interfaces, very slender structures, or sharp concave corners, a global tangential shift may lead to a non-uniform redistribution of the parameter pre-images and may no longer be sufficient to maintain patch quality. In such cases, we expect that the single global shift would have to be replaced by a localized redistribution of the pre-images (for instance weighted by local arc length and curvature, or applied per patch rather than globally), which is left to future work. The minimum scaled Jacobian $J_{\min}$ monitored throughout the simulations provides a direct indicator of whether any such extension preserves bijectivity and mesh quality.

3.3. ALE mesh velocity

Because the patch maps are constructed anew at each time step, the mesh velocity is not defined by incremental node displacement but by the difference between the new and previous patch maps at the same parametric coordinate. For patch k , we define the parameter-space mesh velocity

$$\hat{\mathbf{w}}_k^n(\xi) := \frac{\mathbf{F}_k^n(\xi) - \mathbf{F}_k^{n-1}(\xi)}{\Delta t}, \quad \xi \in \mathcal{P}. \quad (29)$$

In the NURBS representation $\mathbf{F}_k^n(\xi) = \sum_{i=1}^{n_k} N_i(\xi) \mathbf{x}_{k,i}^n$, this reduces to

$$\hat{\mathbf{w}}_k^n(\xi) = \sum_{i=1}^{n_k} N_i(\xi) \frac{\mathbf{x}_{k,i}^n - \mathbf{x}_{k,i}^{n-1}}{\Delta t}, \quad (30)$$

i.e., the NURBS interpolant of the control-point displacements divided by Δt .

The mesh velocity used in the ALE momentum equation (5a) is the push-forward of $\hat{\mathbf{w}}_k^n$ to the current fluid domain $\Omega^f(t^n)$:

$$\mathbf{w}^n(\mathbf{x}) := \hat{\mathbf{w}}_k^n((\mathbf{F}_k^n)^{-1}(\mathbf{x})), \quad \mathbf{x} \in \Omega_k(t^n). \quad (31)$$

In practice, the inverse $(\mathbf{F}_k^n)^{-1}$ is never formed explicitly: (31) is evaluated at quadrature points whose parametric coordinates ξ_q are known by construction, so that $\mathbf{w}^n(\mathbf{F}_k^n(\xi_q)) = \hat{\mathbf{w}}_k^n(\xi_q)$ is computed directly from the control-point increments.

3.4. Solution transfer and the discrete geometric conservation law

Since the geometric parameterizations are constructed at every time step, each fluid solve is posed on a new parameterization $\Omega_h^{f,n}$ that differs from the previous one $\Omega_h^{f,n-1}$. The discrete fluid fields, therefore, have to be transferred across meshes between steps. This transfer problem is distinct from the geometric consistency of a conservative ALE time discretization: the former concerns the representation of the old field on the new parameterization, whereas the latter concerns the discrete evolution of the geometric Jacobian and its consistency with the mesh flux.

Transfer operator. For each patch $k = 1, \dots, n_p$ and each point $\mathbf{x} \in \Omega_k(t^n)$, the composition

$$\Phi_k^n(\mathbf{x}) := \mathbf{F}_k^{n-1}((\mathbf{F}_k^n)^{-1}(\mathbf{x})) \in \Omega_k(t^{n-1}) \quad (32)$$

identifies the point on the previous patch that shares the same parametric coordinate as $\mathbf{x}$. Although (32) is written in physical coordinates, the inverse map $(\mathbf{F}_k^n)^{-1}$ is not formed explicitly. Instead, the correspondence is evaluated directly in parameter space: for a point $\mathbf{x}^n = \mathbf{F}_k^n(\xi)$ on the current parameterization, the corresponding point on the previous parameterization is

$$\Phi_k^n(\mathbf{x}^n) = \mathbf{F}_k^{n-1}(\xi). \quad (33)$$

Thus, no point-wise Newton solve is required to evaluate the same-parametric-coordinate correspondence between consecutive parameterizations.

Given the discrete velocity $\mathbf{u}_h^{f,n-1}$ on $\Omega_h^{f,n-1}$, the transferred field on the new mesh is defined by the quasi-interpolation

$$\tilde{\mathbf{u}}_h^{f,n} := \mathcal{Q}_h^n(\mathbf{u}_h^{f,n-1} \circ \Phi^n), \quad (34)$$

where Φ^n denotes the collection $\{\Phi_k^n\}_{k=1}^{n_p}$ applied patch-wise, and $\mathcal{Q}_h^n$ is a quasi-interpolation operator onto the discrete velocity space on $\Omega_h^{f,n}$. The pressure field is transferred in the same manner with $\mathcal{Q}_h^n$ replaced by its analogue on the pressure space.

Quasi-interpolation operator. For a scalar function f on $\Omega_h^{f,n}$, the quasi-interpolant takes the form

$$\mathcal{Q}_h^n f = \sum_{i=1}^{n_{\text{dof}}} \lambda_i^n(f) N_i^n, \quad (35)$$

where $\{N_i^n\}$ is the global basis on $\Omega_h^{f,n}$ and each λ_i^n is a local linear functional supported in a small neighborhood of the i th control point. In this work we use de Boor–Fix-type functionals [39], for which $\mathcal{Q}_h^n$ is local, linear, and constant-preserving:

$$\mathcal{Q}_h^n c = c \quad \text{for every constant } c \in \mathbb{R}. \quad (36)$$

Because each λ_i^n depends only on the values of f in a small neighborhood, the cost of (35) scales linearly with the number of degrees of freedom and is negligible compared with the cost of the fluid solve. For vector-valued fields, (35) is applied component-wise.

Discrete geometric conservation law. The DGCL [19,40] is commonly expressed through free-stream preservation: a fully discrete ALE scheme should preserve a spatially uniform velocity field under arbitrary mesh motion,

$$\mathbf{u}_h^{f,0} = \mathbf{u}_0 \in \mathbb{R}^d \implies \mathbf{u}_h^{f,n} = \mathbf{u}_0 \quad \forall n \geq 1. \quad (37)$$

Preserving (37) is widely regarded as a minimal consistency requirement for ALE schemes on moving meshes and is often linked to their nonlinear stability [19,20]; see also [22] for a discussion in the IGA context.

Proposition 1. *Let the ALE momentum equation (5a) be discretized in time by backward Euler¹ with mesh velocity $\mathbf{w}^n$ defined by (31). In the absence of external forcing and with boundary data compatible with a constant state, the per-step parameterization and transfer procedure of Sections 3.2 and 3.4 preserves a spatially uniform velocity field in the sense of (37).*

Proof. Suppose $\mathbf{u}_h^{f,n-1} = \mathbf{u}_0$ on $\Omega_h^{f,n-1}$ for some spatially uniform $\mathbf{u}_0 \in \mathbb{R}^d$.

Since $\mathbf{u}_0$ is a constant field, its pullback by any map equals itself; in particular,

$$\mathbf{u}_0 \circ \Phi^n = \mathbf{u}_0 \quad \text{on } \Omega_h^{f,n}.$$

Applying $\mathcal{Q}_h^n$ to both sides and using the constant-preserving property (36) component-wise yields

$$\tilde{\mathbf{u}}_h^{f,n} = \mathcal{Q}_h^n(\mathbf{u}_0 \circ \Phi^n) = \mathcal{Q}_h^n \mathbf{u}_0 = \mathbf{u}_0.$$

For a spatially uniform velocity field $\mathbf{u}_0$, the strain rate and the convective transport in (5a) vanish identically:

$$\mathbb{D}(\mathbf{u}_0) = 0, \quad ((\mathbf{u}_0 - \mathbf{w}^n) \cdot \nabla) \mathbf{u}_0 = \mathbf{0}, \quad \nabla \cdot \mathbf{u}_0 = 0.$$

Consequently, the viscous term vanishes and the continuity equation is trivially satisfied; the pressure field adjusts to balance boundary data and body forces, but does not enter the velocity update. The backward Euler momentum balance therefore reduces to

$$\frac{\rho^f}{\Delta t} (\mathbf{u}_h^{f,n} - \tilde{\mathbf{u}}_h^{f,n}) = \mathbf{0},$$

which gives $\mathbf{u}_h^{f,n} = \tilde{\mathbf{u}}_h^{f,n} = \mathbf{u}_0$. The claim follows by induction on n . $\square$

Fully discrete matrix form. Proposition 1 concerns the advective form of the convective term, which is the form used in all computations reported in Section 4. In that form the constant-preserving transfer is by itself sufficient, because every spatial term vanishes pointwise on a uniform state and the parameterization Jacobian never appears. The discrete geometric conservation law is, however, classically stated for the conservative form, where the Jacobian enters the discrete time derivative explicitly and does not drop out; there a constant-preserving transfer alone does not establish free-stream preservation. We therefore give the corresponding matrix-level statement, both to show that the present construction is DGCL-consistent in this stronger sense and because it is what is required should the geometry module be combined with a conservative ALE discretization. At the fully discrete level, the time-dependent contribution is assembled as

$$\frac{M^{n+1} \mathbf{U}^{n+1} - M^n \tilde{\mathbf{U}}^n}{\Delta t} - C_{\text{GCL}}^{n+\frac{1}{2}} \mathbf{U}^{n+1}, \quad (38)$$

where $\tilde{\mathbf{U}}^n$ is the coefficient vector transferred from the preceding parameterization. The matrices below are defined for a scalar basis and applied separately to each velocity component.

Let

$$\hat{\Omega} := \bigcup_{k=1}^{n_p} \mathcal{P}_k, \quad \mathcal{J}^m|_{\mathcal{P}_k} = J_k^m := \det \mathcal{J}_k^m = \det \left(\frac{\partial \mathbf{F}_k^m}{\partial \xi} \right), \quad m \in \{n, n+1\},$$

¹ The proof extends with essentially no change to any consistent single-step integrator: for a θ -method with $\theta \in (0, 1]$, all spatial terms in (5a)–(5b) vanish pointwise for a constant velocity field, and the velocity update reduces to $\mathbf{u}_h^{f,n} = \tilde{\mathbf{u}}_h^{f,n}$. In Section 4, $\theta = 1$ is used for the perpendicular-flap, rotating-square, and three-dimensional rotor examples, and $\theta = 1/2$ (Crank–Nicolson) for the Turek–Hron FS12 benchmark,

denote the disjoint union of the patch parameter domains and the corresponding parameterization Jacobian determinants. The functions ϕ_i are the reference-domain analysis basis functions, understood patch-wise on $\hat{\Omega}$. We define $C_{\text{GCL}}^{n+\frac{1}{2}}$ as the interval-averaged Jacobian-rate correction matrix over $[t^n, t^{n+1}]$:

$$(M^n)_{ij} := \int_{\hat{\Omega}} J^n \phi_i \phi_j d\xi, \quad m \in \{n, n+1\}, \quad (C_{\text{GCL}}^{n+\frac{1}{2}})_{ij} := \int_{\hat{\Omega}} \frac{J^{n+1} - J^n}{\Delta t} \phi_i \phi_j d\xi. \quad (39)$$

For the linear-in-time parameterization between t^n and t^{n+1} , the metric identity $\partial_t J = J \nabla \cdot \mathbf{w}$ identifies the second integral with the time-averaged mesh-divergence contribution. Its evaluation through the endpoint Jacobians enforces that relation algebraically.

Proposition 2 (Fully Discrete DGCL Identity). *Suppose that M^n , M^{n+1} , and $C_{\text{GCL}}^{n+\frac{1}{2}}$ are assembled with the same analysis basis and quadrature functional on the common knot refinement formed from both consecutive parameterizations. Then*

$$C_{\text{GCL}}^{n+\frac{1}{2}} = \frac{M^{n+1} - M^n}{\Delta t}. \quad (40)$$

Consequently, if the transfer operator reproduces constants as in (36), then (38) vanishes for $\tilde{\mathbf{U}}^n = \mathbf{U}^{n+1} = \mathbf{U}_0$. Hence, the fully discrete update satisfies the DGCL condition (37) for any pair of consecutive admissible parameterizations.

Proof. Let $\mathcal{J}$ denote the numerical quadrature rule shared by the three matrices, using the same integration cells, quadrature points, and weights on the common knot refinement. Applying its linearity entrywise to (39) gives

$$\begin{aligned} (C_{\text{GCL}}^{n+\frac{1}{2}})_{ij} &= \frac{1}{\Delta t} \mathcal{J}[(J^{n+1} - J^n) \phi_i \phi_j] \\ &= \frac{1}{\Delta t} (\mathcal{J}[J^{n+1} \phi_i \phi_j] - \mathcal{J}[J^n \phi_i \phi_j]) = \frac{(M^{n+1})_{ij} - (M^n)_{ij}}{\Delta t}, \end{aligned} \quad (41)$$

which proves (40). Substitution into (38) yields

$$\frac{M^n(\mathbf{U}^{n+1} - \tilde{\mathbf{U}}^n)}{\Delta t}. \quad (42)$$

For a constant free stream $\mathbf{u}_0$, let $\mathbf{U}_0$ denote its coefficient vector. The transfer property gives $\tilde{\mathbf{U}}^n = \mathbf{U}_0$. Setting $\mathbf{U}^{n+1} = \mathbf{U}_0$ therefore makes (42) vanish. Under compatible boundary data and zero forcing, the remaining spatial terms vanish as in Proposition 1. The update therefore leaves $\mathbf{U}_0$ unchanged to machine precision. $\square$

Remark 1. The proof uses the advective form of the convective term,

$$(\mathbf{u}^f - \mathbf{w}) \cdot \nabla \mathbf{u}^f,$$

which vanishes identically for any spatially uniform $\mathbf{u}^f$. In the conservative form

$$\nabla \cdot ((\mathbf{u}^f - \mathbf{w}) \otimes \mathbf{u}^f),$$

a constant $\mathbf{u}^f$ leaves a residual proportional to $\nabla \cdot \mathbf{w}$. This term need not vanish separately; it is balanced by the temporal change of the parameterization Jacobian when the discrete identity (40) is satisfied.

Remark 2. Proposition 2 is algebraic and does not require exact integration of the time-dependent Jacobian fields. It requires the old mass matrix, new mass matrix, and Jacobian-rate correction matrix to use the same quadrature functional; the union integration basis enforces this common partition in the implementation.

The identity (40) should not be read as a tautology, since it holds for the particular discretization (39) of the correction term and not for the discretization one would write down first. Assembling $C_{\text{GCL}}^{n+\frac{1}{2}}$ directly from the metric identity $\partial_t J = J \nabla \cdot \mathbf{w}$, that is, as $\int_{\hat{\Omega}} J^{n+\frac{1}{2}} (\nabla \cdot \mathbf{w}^{n+\frac{1}{2}}) \phi_i \phi_j d\xi$ with the Jacobian and the mesh divergence evaluated at the midpoint, reproduces $(M^{n+1} - M^n)/\Delta t$ only up to a time-interpolation and quadrature error. The free stream would then be preserved to that accuracy rather than to machine precision, and the error would enter at every step. What makes the present construction exact is that the correction is assembled from the two endpoint Jacobians themselves, on the common knot refinement of the two parameterizations and with the same quadrature functional as the two mass matrices, so that the cancellation takes place at the level of the quadrature sums rather than of the integrals they approximate. The requirement is therefore a genuine constraint on the implementation, and it is the reason the union integration basis is used.

Remark 3. The quasi-interpolation operator reproduces constants exactly (36); hence a spatially uniform state is transferred without error. This establishes free-stream preservation for the adopted advective update, but does not make the transfer conservative for integral quantities such as mass or momentum.




Fig. 3. Schematic of one time step of Algorithm 1, illustrated on the rotating-square geometry. **Input:** solution at t^{n-1} on the previous mesh $\Omega_h^{f,n-1}$. **Steps 1–2:** new interface curves obtained from the structural solver or prescribed kinematics, after tangential reparameterization (closed domains). **Step 3:** initial ruled-surface parameterization before quality optimization. **Step 4:** optimized parameterization $\Omega_h^{f,n}$ after the barrier-patch optimization. **Steps 5–7:** the solution on the optimized mesh, obtained by solution transfer, mesh-velocity evaluation, and the fluid solve.

The complete time-stepping procedure of the proposed framework is summarized in Algorithm 1; an illustration on the rotating-square geometry is given in Fig. 3.

Algorithm 1: One time step of the parameterization-driven ALE procedure.

- Input:** $(\mathbf{u}_h^{f,n-1}, p_h^{f,n-1})$ on $\Omega_h^{f,n-1}$
Output: $(\mathbf{u}_h^{f,n}, p_h^{f,n})$ on $\Omega_h^{f,n}$
- 1 Obtain boundary curves $\{\gamma_k^{\text{in}}(\cdot, t^n)\}$ from prescribed kinematics or from the structural solver;
 - 2 **if** closed-domain configuration **then**
 - 3 | Apply the tangential reparameterization (25) to shift the inner boundary;
 - 4 **end**
 - 5 Construct the initial patch mappings $\{\mathbf{F}_k^n\}$ from the boundary curves: ruled surface (24) for closed domains, transfinite interpolation from $\{\gamma_k^{\text{in}}, \gamma_k^{\text{out}}\}$ for open domains (Section 3.2.2);
 - 6 Refine the interior control points by minimizing the barrier-patch functionals (22)–(23);
 - 7 Compute $(\mathbf{u}_h^{f,n}, p_h^{f,n})$ using the transfer operator $\mathcal{Q}_h^n$ in (34);
 - 8 Evaluate the mesh velocity $\mathbf{w}^n$ from (31);
 - 9 Solve (5a)–(5b) on $\Omega_h^{f,n}$ to obtain $(\mathbf{u}_h^{f,n}, p_h^{f,n})$;
-

3.5. Bijectivity criterion

We now specify the discrete condition used to enforce the bijectivity requirement (8) on the patch maps $\{\mathbf{F}_k^n\}$ produced at each time step by the barrier-patch procedure of Section 3.2.1.

The patch-wise Jacobian matrix

$$\mathcal{J}_k^n(\xi) := \frac{\partial \mathbf{F}_k^n}{\partial \xi}(\xi) \quad (43)$$

is the discrete, time-indexed realization of the ALE Jacobian $\mathcal{J}^{\text{ALE}}$ introduced in (7). A standard sufficient condition for local invertibility of $\mathbf{F}_k^n$ is that

$$\det \mathcal{J}_k^n(\xi) > 0 \quad \text{for all } \xi \in \mathcal{P}, k = 1, \dots, n_p. \quad (44)$$

A vanishing determinant signals a local collapse of the parameterization; a sign change corresponds to a fold-over and hence to loss of injectivity.

In practice, (44) is monitored through the minimum scaled Jacobian

$$\mathcal{J}_{\min}(t^n) := \min_{1 \leq k \leq n_p} \min_{\xi \in S} \frac{\det \mathcal{J}_k^n(\xi)}{\prod_{a=1}^d \|\partial_{\xi_a} \mathbf{F}_k^n(\xi)\|}, \quad (45)$$

where $S \subset \mathcal{P}$ is the set of Gauss quadrature points used in the element-wise integration. By Hadamard's inequality, $\mathcal{J}_{\min}(t^n) \in [-1, 1]$, with $\mathcal{J}_{\min}(t^n) = 1$ attained if and only if the columns of $\mathcal{J}_k^n$ are pairwise orthogonal at every sampling point (i.e., the parameterization is locally orthogonal in the ξ_a -directions). A parameterization is considered admissible whenever $\mathcal{J}_{\min}(t^n) > 0$; values close to zero indicate proximity to degeneracy. The indicator $\mathcal{J}_{\min}$ is tracked throughout the numerical experiments of Section 4.

4. Numerical experiments

This section assesses the parameterization-driven ALE framework through six studies spanning kinematic parameterization tests, coupled FSI benchmarks, and three-dimensional large-motion configurations.

The first case (Section 4.1) is purely kinematic: the fluid parameterization is driven by prescribed rigid-body rotation of an embedded square, with no flow solver involved. This isolates the parameterization itself and quantifies the range of cumulative rotation angles for which the barrier-patch construction remains valid. This test has two purposes. First, it provides a stress test for the parameterization method without tangential slip. Second, the comparison between the higher-order spline result and the first-order case $p = 1$ shows that the higher degree and inter-element continuity brought by the spline representation help preserve parameterization validity over a larger admissible rotation range under large accumulated deformation.

The perpendicular-flap benchmark from the preCICE tutorial suite (Section 4.2) verifies the response in a moderate-deformation two-way coupled setting, while the Turek–Hron FSI2 benchmark (Section 4.3) tests sustained large-amplitude periodic deformation. The sphere-in-channel study (Section 4.4) extends the assessment to three-dimensional large translation: its prescribed-motion part compares mesh quality with Di Cristofaro et al. [31], whereas its free-fall part validates the coupled velocity response against published data [41,42]. The flow past a rotating square (Section 4.5) then returns to the geometry of the first case under continuous large-angle rotation and uses both isogeometric and finite-element flow solvers to demonstrate solver portability. Finally, the three-dimensional rotor tests the volumetric construction under sustained rotation. Thus, the perpendicular-flap, Turek–Hron, and free-fall cases provide dynamical comparisons with external data, while the prescribed sphere, rotating square, and rotor isolate robustness under large translation or rotation.

The fluid is discretized with Taylor–Hood-type isogeometric spaces (Section 3.1) on the multi-patch geometries constructed in Section 3.2. All fluid solves are performed in G+Smo [43,44]; the rotating-square case additionally uses DEAL.II [45] as an external finite-element solver. The three two-way coupled benchmarks (perpendicular flap, Turek–Hron FSI2, and free fall) are coupled in a partitioned manner through the preCICE library [46].

4.1. Admissible rotation range of the mesh parameterization

Our first example isolates the parameterization itself, independent of any flow or FSI coupling, and quantifies the range of cumulative rotation angles for which it remains valid. We consider a square of side length L centered in a $3L \times 3L$ outer box; the annular region is partitioned into four body-fitted patches, matching the inner block of the fluid domain used in Section 4.5. The square is rotated incrementally about its center, and each mesh-update method is applied step by step until either $J_{\min}$ drops to zero (parameterization failure) or a target cumulative rotation is reached. The indicator of interest is $J_{\min}(\theta)$ from (45), tracked as a function of the rotation angle θ .

Maximum sustainable rotation angle. We compare the proposed barrier-patch (BP) parameterization with six classical mesh-update strategies: harmonic extension (HE), incremental harmonic extension (IHE), linear elasticity (LE), incremental linear elasticity (ILE), tangential incremental nonlinear elasticity (TINE), and bi-harmonic extension (BHE); detailed formulations are given in [49]. The maximum rotation angle sustained by each method is reported in Fig. 4. All six classical methods lose validity well before 90° ; the best performer among them, bi-harmonic extension, fails at 67.1° . The Barrier Patch Method remains valid up to 127.5° .

The full benefit of the proposed approach relies on the spline parameterization, and in particular on its degree and inter-element continuity, which together yield smoother Jacobian fields; these are essential when the domain undergoes large accumulated deformation. To quantify this effect, we repeated the stress test with a first-order ($p = 1$) parameterization, which lacks the inter-element smoothness the method is designed to exploit, while keeping the regeneration algorithm, the patch layout, and the refinement level unchanged. As shown in Fig. 5, the maximum admissible rotation angle drops to 31.35° , far below the 127.5° reached by the higher-order barrier-patch parameterization in Fig. 4. Higher-order spline parameterizations thus provide a substantially larger resistance to large mesh deformations; all subsequent examples therefore employ parameterizations of degree $p \geq 2$.

4.2. Perpendicular flap

As the first coupled configuration, we consider the perpendicular-flap benchmark from the preCICE tutorial suite [46]. The benchmark models two-dimensional fluid–structure interaction between channel flow and an elastic flap mounted perpendicular to the bottom wall; the incoming flow exerts pressure on the flap and induces oscillatory motion. Reference solutions from several solver combinations are available, and the accuracy of our isogeometric solver and the spline-based coupling interface on this benchmark has been established in [12]. Here we therefore focus on the parameterization component and compare it to classical mesh-update strategies.

Problem setup. The physical domain, geometric dimensions, and material parameters are summarized in Fig. 6. The flow is incompressible; a uniform horizontal velocity is prescribed on Γ_{inflow} , a free boundary condition on Γ_{outflow} , and no-slip on the top and bottom walls. The flap is clamped on Γ_{fixed} and loaded by the fluid traction on its wetted boundary Γ_{fs} . The structure is modeled as a St. Venant–Kirchhoff solid. Time integration is carried out with the backward Euler scheme up to $T = 5$ s using a constant time step $\Delta t = 0.01$ s. The two subproblems are coupled in a partitioned manner through preCICE using an implicit coupling scheme with IQN-ILS quasi-Newton acceleration [50].




Fig. 4. Maximum admissible rotation angles for different mesh-motion strategies. Each panel in the left block shows the parameterization at the maximum angle the corresponding classical method can handle before losing validity, with the angle indicated in the upper-right chip. The right column shows the parameterization-driven Barrier Patch Method at three rotation angles (45° , 90° , and its admissible limit 127.5°), demonstrating that the parameterization remains valid throughout the rotation and reaches a substantially larger maximum angle than any classical scheme (see [47,48]).

Snapshot of the coupled response. Fig. 7 shows the fluid velocity magnitude together with the regenerated isogeometric parameterization at six instants spanning the simulation interval. As the flap oscillates in response to the channel flow, the surrounding patches deform substantially near the flap tip, yet the barrier-patch parameterization keeps every element well-shaped at every time step.

Comparison with classical mesh-update strategies. We compare the barrier-patch (BP) parameterization with five classical mesh-update strategies, while keeping the fluid and structural discretizations identical: harmonic extension (HE), incremental harmonic extension (IHE), linear elasticity (LE), incremental linear elasticity (ILE), and tangential incremental nonlinear elasticity (TINE).

Fig. 8 shows the tip-displacement histories over $T = 5$ s. The six curves are visually indistinguishable on the plotted scale, so the barrier-patch parameterization reproduces the response of the established mesh-update strategies in this moderate-deformation regime.

Fig. 9 reports the minimum scaled Jacobian of the fluid ALE mesh over $[0, 5]$ s. All six methods remain strictly positive throughout, confirming the validity of each parameterization. HE, IHE, LE, ILE, and TINE track the oscillation of the flap: $J_{\min}$ decreases during phases of large tip displacement and recovers as the flap returns toward its undeformed configuration. The barrier-patch method, by contrast, maintains $J_{\min}$ above 0.6 with only minor fluctuations, yielding a mesh quality that is effectively decoupled from the structural motion.

4.3. Turek-Hron FSI 2

The second benchmark is the Turek–Hron FSI2 problem [51], a standard test case for fluid–structure interaction. It models the two-dimensional flow-induced oscillations of an elastic beam clamped to the downstream side of a rigid cylinder. The parameters





Fig. 5. Maximum admissible rotation angles for first-order ($p = 1$) parameterizations with 2 uniform refinements, using a 0.05° angle resolution and without inner-boundary tangential slip. Each panel in the left block shows the parameterization at the maximum sampled angle the corresponding classical method can handle before losing validity. The right column shows the Barrier Patch Method at two intermediate angles (20° and 30°) and its admissible limit 31.35° .



DIMENSIONS		
L	6.00	m
H	4.00	m
l	2.95	m
h	1.00	m
w	0.10	m
FLUID		
ρ^f	1.0	kg/m ³
ν^f	1.0×10^{-5}	m ² /s
SOLID		
E	4×10^6	N/m ²
ν^s	0.3	–
ρ^s	3×10^3	kg/m ³

Fig. 6. Setup and parameters of the perpendicular-flap benchmark. The geometry schematic is adapted from the preCICE documentation [46].




Fig. 7. Time history of the fluid velocity magnitude and the isogeometric parameterization for the perpendicular-flap benchmark, shown at six instants spanning $[0.1, 5.0]$ s. The flap oscillates in response to the channel flow, and the surrounding patches are reparameterized at each time step to remain valid throughout the deformation. All panels share the common color bar shown above each frame.




Fig. 8. Perpendicular-flap benchmark: horizontal tip displacement histories obtained with six mesh-update strategies.

of the FSI2 configuration place the system in a regime of large-amplitude self-sustained oscillations, making it a demanding test of both solver accuracy and mesh robustness under sustained deformation.

Problem setup. The physical domain, shown in Fig. 10, is a rectangular channel of length $L = 2.5$ m and height $H = 0.41$ m. A rigid circular cylinder of diameter $D = 0.1$ m is centered at $(0.2, 0.2)$ m, with an elastic beam of length $l = 0.35$ m and thickness $h = 0.02$ m clamped to its downstream side.

A parabolic velocity profile is prescribed at the inflow boundary Γ_{in} ,

$$v_{\text{in}}(y) = 1.5 \bar{v} \frac{4y(H-y)}{H^2}, \quad (46)$$

with mean velocity $\bar{v} = 1$ m/s, corresponding to a Reynolds number $\text{Re} = 100$. No-slip conditions are imposed on the upper and lower walls, and a stress-free condition is applied at the outlet Γ_{out} .




Fig. 9. Perpendicular-flap benchmark: minimum scaled Jacobian $J_{\min}$ of the fluid ALE mesh over time. The barrier-patch method maintains $J_{\min}$ in a narrow band around 0.70.




Fig. 10. Turek-Hron FSI2 benchmark setup: a rigid cylinder with an elastic beam clamped to its downstream side, immersed in a channel flow. The inflow is a parabolic profile; the outflow boundary is stress-free.



Fig. 11. Patch decomposition of the fluid domain for the FSI2 benchmark. The rigid cylinder (white disc) and the elastic beam (white strip) form the inner obstacle. Three colored patches conforming to the cylinder-beam interface (purple, magenta, and green) constitute the *regenerated* region: at every coupling iteration these patches are reparameterized to follow the deformed beam, using the barrier-patch construction of Section 3.2. The remaining gray region is covered by *frozen* extension patches that span the rest of the channel up to the inflow, outflow, and wall boundaries; these are generated once and remain fixed throughout the simulation. Confining the regeneration to a narrow neighborhood of the beam keeps the per-iteration meshing cost essentially independent of the channel length. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

The fluid has density $\rho^f = 10^3 \text{ kg/m}^3$ and dynamic viscosity $\mu^f = 1.0 \text{ Pa} \cdot \text{s}$. The solid is modeled as a St. Venant-Kirchhoff material with density $\rho^s = 10^4 \text{ kg/m}^3$, shear modulus $\mu^s = 0.5 \times 10^6 \text{ Pa}$, and Poisson's ratio $\nu^s = 0.4$. Time integration uses a constant step $\Delta t = 0.0025 \text{ s}$ over the interval $[0, 18] \text{ s}$.

Fluid-structure coupling follows the scheme used in the previous benchmark: a partitioned implicit solver with Aitken relaxation [52]. The fluid parameterization is handled by the barrier-patch parameterization of Section 3.2; at each coupling iteration, only the patches adjacent to the elastic beam are regenerated, while the surrounding extension patches remain fixed.

Patch layout. The fluid domain is decomposed into a small set of spline patches assembled around the cylinder-beam obstacle, as illustrated in Fig. 11. Three patches form a thin layer that conforms to the obstacle: two wedge-shaped patches above and below the beam, and a patch closing off its tip. These three patches are regenerated at every coupling iteration so that their boundaries track the deformed beam exactly. The remainder of the channel (the inflow region upstream of the cylinder, the long downstream wake, and the upper and lower wall strips) is tiled by a fixed set of extension patches that are constructed once and never modified.




Fig. 12. Turek-Hron FSI2 benchmark: velocity field and regenerated mesh at two instants in the periodic regime, near opposite extrema of the beam oscillation. The velocity panels share the common color bar shown at the top.

during the simulation. This choice isolates the geometric variability to a narrow neighborhood of the beam, where the deformation is concentrated, and keeps the cost of each remeshing operation independent of the channel length.

Results. We evaluate the barrier-patch parameterization along three dimensions: the validity and quality of the regenerated mesh under sustained large deformation, the accuracy of the FSI response against published reference values, and the computational overhead introduced by parameterization within the coupling loop.

Fig. 12 presents the velocity field and the regenerated parameterization at two instants in the periodic regime, chosen near opposite extrema of the beam oscillation. The barrier-patch parameterization delivers a well-shaped parameterization in both configurations, with no visible element degradation even at maximum tip deflection. Across the full simulation window, the minimum scaled Jacobian remains strictly positive, confirming that parameterization validity is preserved throughout the sustained large-amplitude motion.

To assess solution accuracy, Fig. 13 reports the time history of the vertical tip displacement of the elastic beam. After an initial transient of roughly 5 s, the response settles into a periodic regime whose amplitude agrees closely with the reference value ± 80.6 mm reported by Turek and Hron [51], indicating that the barrier-patch parameterization does not compromise the fidelity of the coupled FSI response.

Fig. 14 reports the cumulative wall-clock time, decomposed into contributions from the beam solver, the flow solver, and the barrier-patch parameterization. The flow solve dominates the total cost, while parameterization contributes a modest share that grows approximately linearly with simulation time. The barrier-patch procedure therefore imposes an overhead comparable to that of conventional mesh-motion schemes, despite offering the stronger guarantee of strictly positive Jacobians under sustained large deformation.

4.4. Sphere-in-channel benchmarks in three dimensions

To assess the method under large translational motion, we consider two three-dimensional sphere-in-channel configurations. The prescribed-translation benchmark of Di Cristofaro et al. [31] tests whether a fixed spline space and patch topology remain valid over a 40 m displacement and permits direct comparison of failure time, mesh reconnection, and mesh-size growth. The free-fall benchmark then tests the two-way coupled response against published terminal-velocity data.

4.4.1. Prescribed large translation

Problem setup. We first consider the prescribed-translation case of Di Cristofaro et al. [31], illustrated in Fig. 15. The fluid occupies a cylinder of length $L = 60$ m and diameter $D_2 = 20$ m, and contains a concentric rigid sphere of diameter $D_1 = 8$ m. Thus the channel and sphere radii are 10 m and 4 m, respectively. With the axial coordinate spanning $0 \leq z \leq 60$ m, the sphere center starts at $z_c(0) = 50$ m above the lower base and is assigned the velocity

$$\mathbf{u}^f|_{\Gamma_s} = -5 \mathbf{e}_z \text{ m/s.} \quad (47)$$




Fig. 13. FSI2 benchmark: time history of the vertical tip displacement of the elastic beam. The dashed orange lines mark the reference amplitude ± 80.6 mm reported by Turek and Hron [51].




Fig. 14. FSI2 benchmark: cumulative wall-clock time decomposed into the beam solver, the flow solver, and the barrier-patch mesh-regeneration step.

Homogeneous Dirichlet conditions are imposed on the remaining fixed boundary. As in the reference study, $\rho^f = 1 \text{ kg/m}^3$ and $\mu^f = 1 \text{ Pa}\cdot\text{s}$. The interval $[0, 8] \text{ s}$ is discretized with $\Delta t = 0.02 \text{ s}$, during which the sphere translates by 40 m. The fluid domain is represented by six volumetric spline patches, in the same O-grid topology used for the free-fall case of Section 4.4.2, and is reconstructed independently from the current sphere position at each step; no connectivity change or refinement is performed, so the element and control-point counts are identical at every one of the 400 steps. The geometry and pressure space use degree $(2, 2, 2)$, while the velocity space is obtained by one degree elevation and uses degree $(3, 3, 3)$, giving a Taylor–Hood Q_3/Q_2 pair. The discretization contains $3 \times 3 \times 2 = 18$ volume elements per patch and 108 volume elements in total. After identification of conforming patch-interface degrees of freedom and elimination of Dirichlet unknowns, the coupled system contains 4052 global free degrees of freedom, comprising 3660 velocity degrees of freedom (1220 per component) and 392 pressure degrees of freedom.

Results. Fig. 16 shows the fluid response alongside the regenerated volumetric parameterization at the two latest comparison times.

Panels (c,d) are colored by the scaled Jacobian J_s , i.e., the quantity minimized in (45). Since J_s is normalized by the lengths of the parametric tangent vectors, it reflects their relative orientation rather than the cell aspect ratio. The long cells in the upper central patches are nearly orthogonal and therefore have values close to one; the more skewed cells in the neighboring patches have lower values. Jumps at patch interfaces are expected because the geometry is C^0 -continuous there, whereas its patch-wise derivatives need not be.

$J_{\min}$ remains strictly positive at every sampled configuration, confirming that the discrete bijectivity condition (44) holds throughout the translation. Di Cristofaro et al. [31], by contrast, report that their tetrahedral mesh fails at $t = 2.0 \text{ s}$ without reconnection, and that continuing the run requires reconnection together with an approximately 10% increase in nodes and elements; the present six-patch parameterization completes the full 8.0 s run on the same fixed connectivity, with no growth in the number of




Fig. 15. Prescribed-translation setup based on Di Cristofaro et al. [31]. Panel (a) shows the sphere moving vertically from $z_c = 50$ m to $z_c = 10$ m with velocity $-5 \mathbf{e}_z$ m/s. Panel (b) shows the transverse section through the sphere center, with channel diameter $D_2 = 20$ m and sphere diameter $D_1 = 8$ m.




Fig. 16. Center-plane views of the prescribed-translation test at $t = 6.6$ s and 8.0 s, selected to match the two snapshot times reported by Di Cristofaro et al. [31]. Panels (a,b) show the fluid-velocity magnitude with the common range $0 \leq \|\mathbf{u}'\| \leq 5.2$ m/s. Panels (c,d) show the regenerated volumetric parameterization, colored by the cellwise scaled Jacobian J_s on the common rainbow scale $[0, 1]$; the overlaid lines show the parameterization grid. The white disk denotes the rigid sphere. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)




Fig. 17. Minimum scaled Jacobian $J_{\min}$ over all spline patches during the prescribed translation, evaluated on 85 sampled configurations of the exported volumetric mesh. The global minimum and maximum are marked in red and green, respectively. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

control points or elements. As Fig. 17 shows, $J_{\min}$ stays within the narrow band $[0.0722, 0.0932]$ throughout, reaching its minimum at the two symmetric endpoints $t = 0$ and $t = 8$ s — where the sphere sits only 10 m from the nearest end wall — and its maximum at $t = 4$ s, when the sphere is equidistant from both ends.

4.4.2. Two-way coupled free fall

Problem setup. We next consider the freely falling sphere introduced by Casquero et al [41] and revisited by Black et al. [42], whose finest-mesh solution we use as the reference in the comparison below. As in Section 4.4.1, the fluid domain is discretized into six volumetric spline patches conforming to a concentric sphere-in-cylinder geometry, here at a much smaller scale and with the sphere free to translate under the coupled fluid and gravitational forces rather than prescribed. A cylinder of radius $A = 2$ cm and height 4 cm contains a sphere of radius $a = 0.25$ cm, initially centered at $z_c = 2.5$ cm. The fluid parameters are $\rho^f = 1$ g/cm³ and $\mu^f = 10$ dyn s/cm², the sphere density is $\rho^s = 1.5$ g/cm³, and gravity is $\mathbf{g} = -981 \mathbf{e}_z$ cm/s². A traction-free condition is imposed at the top of the cylinder and no slip on the remaining fixed walls. The sphere is treated as rigid and constrained to axial translation, consistently with the analytical terminal-velocity model. Fluid traction and sphere displacement are exchanged by the implicit partitioned coupling scheme through preCICE. We use $\Delta t = 10^{-4}$ s and simulate until $T = 0.1$ s.

Assuming Stokes flow, the wall-corrected terminal speed is

$$v_{\infty} = \frac{1}{K} \frac{2ga^2(\rho^s - \rho^f)}{9\mu^f}, \quad K = \left[1 - 2.10443 \frac{a}{A} + 2.08877 \left(\frac{a}{A} \right)^3 \right]^{-1}, \quad (48)$$

which gives $v_{\infty} = 0.504824$ cm/s. The six-patch fluid parameterization contains $3 \times 3 \times 2 = 18$ volume elements per patch, hence 108 volume elements in total. The fluid is discretized in the same manner as in the prescribed-translation benchmark, using the six-patch Taylor–Hood Q_3/Q_2 spline discretization described in Section 4.4.1.

Results. Fig. 18 shows the fluid velocity magnitude in the sphere’s center plane shortly after release and near the end of the simulated interval, namely at $t = 0.01$ s and $t = 0.10$ s.

Fig. 19 compares the sphere velocity obtained with the present formulation against the reference data.

As shown in Fig. 19, the sphere approaches a terminal speed of 0.5063 cm/s. The relative difference from the wall-corrected value is 0.29%, while the difference from the final value of the finest reference curve is approximately 3.4%. For comparison, Black et al. report a 3.02% difference between their finest-grid result and the analytical value. Throughout the 1000 coupled time steps, the minimum scaled Jacobian $J_{\min}$ of the volumetric parameterization remains in the narrow interval $[0.138429, 0.138617]$. This second test therefore demonstrates that the present formulation yields accurate results under two-way coupling, not only mesh validity: the prescribed test isolates parameterization robustness over a long trajectory, while the free-fall test confirms that the same per-step parameterization remains both stable and accurate when the trajectory is determined by the coupled fluid force.

4.5. Flow past a rotating square

In the kinematic study of Section 4.1 we mapped out the admissible rotation range of the mesh-update strategies. We now return to the same geometry, but couple it to the incompressible Navier–Stokes equations. To validate the present method against an independent reference, we reproduce the rotating-square benchmark of Schalk et al. [30], whose block-structured meshing scheme handles the same connectivity-changing rotation problem in a finite element setting. The benchmark parameters and reference kinetic-energy and momentum histories were obtained directly from the authors and are used here for comparison. The simulation covers one full revolution of the inner square, during which the tangential-slip reparameterization of Section 3.2.3 remains active throughout.




Fig. 18. Two-way coupled free-fall test: fluid-velocity magnitude in the sphere's center plane at $t = 0.01$ s, shortly after release, and at $t = 0.10$ s, close to the terminal speed. The white disk denotes the rigid sphere.




Fig. 19. Two-way coupled free-fall test: downward sphere speed obtained with the present IGA-ALE/preCICE formulation, the finest mesh level reported by Black et al. [42], and the wall-corrected analytical terminal speed from (48). The present discretization uses 108 volume elements with a Taylor-Hood Q_3/Q_2 pair and 4052 global free degrees of freedom (3660 velocity and 392 pressure).

Problem setup. We consider two-dimensional incompressible flow past a rigid square obstacle of side length 2 m, rotating about its center $\mathbf{x}_c = (L/3, H/2)$ with prescribed angular velocity $\dot{\theta} = 2^\circ/\text{s}$. The physical domain, shown in Fig. 20, is a channel of length $L = 25$ m and height $H = 10$ m; the region containing the square is decomposed into four body-fitted patches flush with the inlet and centered on the square, joined to a fixed downstream outflow extension.

A parabolic inflow profile with centerline velocity $u_{\text{in}} = 2$ m/s is prescribed at the inlet. No-slip conditions are imposed on the upper and lower walls, and a stress-free condition is applied at the outlet. On the square boundary, the velocity follows the rigid-body rotation

$$\mathbf{v} = \dot{\theta} \mathbf{e}_z \times (\mathbf{x} - \mathbf{x}_c) = (-\dot{\theta}(y - y_c), \dot{\theta}(x - x_c))^T. \quad (49)$$

The fluid density is $\rho^f = 1.0 \text{ kg/m}^3$ and the kinematic viscosity is $\nu^f = 0.05 \text{ m}^2/\text{s}$.




Fig. 20. Rotating-square benchmark setup, following Schalk et al. [30]: a channel of length $L = 25 \text{ m}$ and height $H = 10 \text{ m}$ contains a rigid square of side 2 m centered at $\mathbf{x}_c = (L/3, H/2)$. The region containing the square is decomposed into four body-fitted patches $\Omega_1, \dots, \Omega_4$, flush with the inlet and centered on the square, and joins a fixed downstream outflow extension Ω_5 .

The fluid domain is decomposed into five patches: four body-fitted patches surrounding the square and one fixed outflow patch. Each patch is parameterized by the barrier-patch method; the spline space is degree-elevated once and uniformly refined four times, with one additional refinement in the obstacle-normal direction on the four near-body patches, as detailed below. This yields 4864 elements and 49,543 mixed velocity–pressure degrees of freedom for a Taylor–Hood Q_4/Q_3 spline pair in the present computations, i.e. tensor-product degree (4, 4) for velocity and (3, 3) for pressure. For reference, Schalk et al. [30] report 11,136 quadrilateral elements with a Taylor–Hood Q_2/Q_1 discretization. Since their global DoF count is not explicitly listed, we report a rough estimate based on standard structured-mesh scaling for Q_2/Q_1 in 2D, DoFs $\approx 9N_e$, which gives about $9 \times 11,136 \approx 1.00 \times 10^5$ mixed velocity–pressure DoFs. We solve in G+Smo [43,44,53] with $\Delta t = 0.005 \text{ s}$ over one revolution ($\theta \in [0^\circ, 360^\circ]$, $T = 180 \text{ s}$) (see Fig. 21).

Validation. Fig. 22 compares the present solution against the reference kinetic-energy and momentum histories of Schalk et al. [30]. Fig. 22(a) shows the kinetic energy E_k of the flow over a full revolution of the square; the two curves overlap closely throughout, including the transient during the first quarter-revolution and the subsequent periodic regime. Fig. 22(b) zooms into the vertical linear momentum P_y over $\theta \in [45^\circ, 160^\circ]$. Both curves show the same beat pattern, in which the oscillation amplitude grows and decays over successive cycles, while remaining in phase. Fig. 22(c) reports the frequency $f = \dot{\theta}/(2\bar{\Delta\theta}_{\text{sw}})$, where $\bar{\Delta\theta}_{\text{sw}}$ is the mean angular interval between consecutive sign changes of P_y on $\theta \in [45^\circ, 180^\circ]$, the start-up transient of the first quarter-revolution being excluded from the average, and the peak magnitude $A = \max_{\theta \in [0^\circ, 360^\circ]} |P_y(\theta)|$. The two frequencies agree to within 0.43% and the amplitudes to within 0.94%; the frequency estimate varies by less than 0.01 percentage points when the start of the averaging window is moved between 40° and 60° . Quantitatively, the kinetic-energy histories of Fig. 22(a) agree to 0.18% in the mean and 0.25% in relative L^2 over the full revolution, with a maximum pointwise deviation of 0.96%. The residual differences are consistent with the small phase and magnitude offsets visible in Fig. 22(b).

Scaled-Jacobian history under sustained rotation. We now return to the diagnostic deferred from Section 4.1 and evaluate $J_{\min}$ of (45) on the fluid parameterization throughout a full revolution of the inner square. Fig. 23 traces the resulting evolution against the rotation angle θ .

Two observations stand out. First, $J_{\min}$ stays strictly bounded away from zero throughout the entire rotation, so no mesh degeneration is encountered. The minimum value $J_{\min} = 0.6085$ is attained at $\theta = 94.5^\circ$, and the maximum $J_{\min} = 0.7048$ at $\theta = 42.75^\circ$; both lie well above the invalidity threshold $J_{\min} = 0$, confirming that the regeneration algorithm satisfies the bijectivity condition of (45). Second, the trajectory exhibits a clear periodic pattern synchronized with the rotation, reflecting the 90° rotational symmetry of the four-patch inner decomposition. The four quarter-revolution cycles are in fact indistinguishable to four significant digits, each attaining the same extrema $J_{\min} = 0.6085$ and $J_{\min} = 0.7048$, so mesh quality does not drift as the rotation accumulates. This is expected rather than fortuitous: each parameterization is rebuilt from the current square orientation alone, so configurations that differ by a multiple of 90° produce the same parameterization up to the symmetry of the patch layout, irrespective of how many revolutions preceded them.




Fig. 21. Velocity field (left) and regenerated ALE mesh (right) at selected rotation angles. The two velocity columns share the common color bar shown at the top. The tangential-slip strategy sustains valid and well-shaped patches far beyond the rotation limit of the plain barrier-patch method.

Cost of tangential slip. A practical question is whether rebuilding the spline parameterization at every time step introduces a significant computational overhead. To quantify this cost, we measure the cumulative mesh-only regeneration time over 200 evenly spaced configurations across eight successive uniform refinement levels ($r = 0, 1, \dots, 7$) and for polynomial degrees $p = 2$ and $p = 3$. For each refinement level, the $p = 2$ and $p = 3$ cases use the same multi-patch layout and the same element topology. Only the spline degree of the parameterization is changed. The reported time includes tangential slip, Barrier Patch, degree elevation, uniform refinement, and topology construction. A higher spline degree increases both the number of quadrature points and the number of active basis functions involved in each local evaluation, even when the number of elements is unchanged.

Fig. 24 summarizes the results. For coarse to moderately refined meshes ($r \leq 4$), the cumulative mesh-only cost remains between 1.92 s and 2.41 s for $p = 2$, and between 1.96 s and 2.43 s for $p = 3$. Beyond this range, the cost grows gradually: at $r = 6$, it reaches 5.38 s for $p = 2$ and 6.27 s for $p = 3$, and at the finest level $r = 7$, it rises to 14.95 s and 21.20 s, respectively.




Fig. 22. Validation against the reference solution of Schalk et al. [30] for the rotating-square benchmark. (a) Kinetic energy E_k of the flow over one full revolution. (b) Zoom in of vertical momentum P_y over $\theta \in [45^\circ, 160^\circ]$. (c) Corresponding frequency and peak magnitude of P_y .




Fig. 23. Minimum scaled Jacobian $J_{\min}$ of the ALE fluid parameterization over one full revolution of the inner square. The red and green circles mark the global minimum and maximum, respectively. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

The levels r in Fig. 24 are obtained by uniformly refining all patches of the layout and therefore form a scaling study rather than a sequence of production runs. The comparison reported above in Fig. 22 instead uses a directionally refined configuration. After four uniform refinements, the four near-body patches receive one additional refinement in the obstacle-normal direction, while no additional tangential refinement is applied. The base discretization comprises one element per near-body patch and eleven elements in the outflow patch, fifteen in total, which is the $r = 0$ entry of Fig. 24. Consequently, each near-body patch contains $2^4 \times 2^5$ elements and the outflow patch contains $11 \times 2^4 \times 2^4$ elements. The total number of elements is therefore

$$N = 4 (2^4 \times 2^5) + 11 (2^4 \times 2^4) = 2048 + 2816 = 4864.$$

This configuration lies between the uniformly refined levels $r = 4$ and $r = 5$ of the sweep. The two finest levels are included primarily as stress tests of the regeneration algorithm: in the IGA setting, the high-order and high-continuity B-spline/NURBS basis provides




Fig. 24. Cumulative mesh-only regeneration time over 200 evenly spaced configurations of the final rotating-square setup, for polynomial degrees $p = 2$ and $p = 3$ at successive uniform refinement levels $r = 0, \dots, 7$. Barrier Patch and tangential slip are enabled; the fluid solve is excluded. The dashed red line marks 2 s and is drawn for reference only. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

Table 1

Cost breakdown of one parameterization-driven ALE time step. Measured times are per-step wall-clock averages over 200 time steps on the production mesh (4864 elements, $\Delta t = 0.005$ s) with 8 OpenMP threads, for polynomial degrees $p = 2$ and $p = 3$. The barrier-patch optimization runs on the pre-refinement single-patch geometry, so its cost is essentially independent of r .

Step in Algorithm 1	Relative cost	Measured (s/step)		Comment
		$p = 2$	$p = 3$	
Tangential reparameterization (line 3)	Negligible	$<10^{-4}$	$<10^{-4}$	One-dimensional curve shift
Initial ruled/transfinite construction (line 5)	Negligible	$<10^{-4}$	$<10^{-4}$	Explicit geometric construction
Barrier-patch optimization (line 6)	Small	0.0075	0.0087	Nonlinear optimization with quadrature evaluations
Quasi-interpolation transfer (line 7)	Dominant geometry cost	0.018	0.020	Local operator, linear in degrees of freedom
Mesh-velocity evaluation (line 8)	Negligible	1.1×10^{-4}	1.2×10^{-4}	Control-point differences
Fluid solve (line 9)	Dominant total cost	1.57	4.04	Navier–Stokes solve (Picard iterations)

a higher approximation capability per element than low-order C^0 finite elements, so resolutions of this order are well beyond what the present benchmark requires. Across all levels considered, the mesh-only regeneration cost per configuration stays at or below approximately 0.11 s, so the parameterization procedure remains computationally affordable relative to the full flow and structural solves (see Table 1).

The measurements confirm the qualitative picture. On the production mesh with 4864 elements, the barrier-patch optimization takes approximately 7.5 ms and 8.7 ms per time step for $p = 2$ and $p = 3$, respectively, while the tangential reparameterization remains below 0.1 ms and the mesh-velocity evaluation costs only about 0.11–0.12 ms. The barrier-patch cost remains relatively insensitive to mesh refinement, since the optimization is performed on the pre-refinement single-patch geometry rather than on the refined computational mesh. The quasi-interpolation transfer requires approximately 18 ms for $p = 2$ and 20 ms for $p = 3$, making it the dominant contribution to the geometry-side cost. Nevertheless, this overhead remains small compared with the fluid solve, which requires approximately 1.57 s and 4.04 s per time step for $p = 2$ and $p = 3$, respectively. Thus, the additional cost introduced by the parameterization-driven ALE procedure constitutes only a small fraction of the overall computational cost.

Portability to a finite element solver. The proposed G+Smo library is stand-alone: it constructs the spline geometry of the physical domain independently of any field solver. Through preCICE, this geometry is passed to DEAL.II [45]; after receiving the spline representation, DEAL.II carries out an independent finite-element flow calculation. The DEAL.II branch therefore serves only as an interoperability test, not as a component or dependency of the geometry library. This workflow is illustrated in Fig. 25.

For the optional portability test, Fig. 26 compares the vertical linear-momentum histories obtained with G+Smo (IGA) and the downstream DEAL.II consumer (FEM) on the same sequence of physical geometries over $\theta \in [45^\circ, 160^\circ]$, matching the angular interval used for the comparison with Schalk et al. in Fig. 22(b). On this interval the two computations agree to 0.51% in the frequency of P_y and to 2.5% in its peak magnitude, with a phase offset that grows monotonically to 0.54° of rotation, roughly one twelfth of a period, by the end of the interval; the pointwise difference is dominated by this drift rather than by a discrepancy in amplitude. A slow relative phase drift is expected here, since the two solvers share only the geometry and differ in basis, quadrature and linear algebra. Combining this with the comparison against Fig. 22(b), and averaging all three over the common interval $\theta \in [45^\circ, 160^\circ]$ covered by the DEAL.II run—hence over a shorter window than the one used in Fig. 22(c)—the frequency of P_y is 0.1472 Hz for Schalk et al. [30], 0.1472 Hz for DEAL.II and 0.1465 Hz for the present isogeometric solver, a spread of 0.51% across three mutually independent flow discretizations driven by the same sequence of regenerated geometries. This confirms that the geometry produced

Geometry library (G+Smo)

Optional consumer (deal.II)




Fig. 25. Stand-alone geometry library and optional finite-element consumer. Through preCICE, G+Smo can operate as an independent geometry library and communicate its spline geometry to a downstream finite-element solver, which constructs its own computational mesh and discrete spaces.




Fig. 26. Comparison of the vertical linear momentum P_y for flow past the rotating square over $\theta \in [45^\circ, 160^\circ]$, computed with G+Smo (IGA) and DEAL.II (FEM) on the same sequence of physical geometries. The frequencies agree to 0.51% and the peak magnitudes to 2.5%, with a phase offset accumulating to 0.54° over the interval. The agreement confirms the portability of the stand-alone G+Smo geometry library to an independent FEM solver.

by the stand-alone library is portable across IGA and FEM solvers. We emphasize that DEAL.II is an optional validation consumer, not a dependency of the proposed library.

5. Application to 3D rotating-rotor benchmark

To demonstrate that the proposed method extends naturally to three dimensions, we apply the parameterization-driven strategy to axial flow through a rotating rotor enclosed in a cylindrical casing.

Problem setup. The physical domain, shown in Fig. 27, is the annular region between a fixed cylindrical outer wall of radius $R = 1.5$ m and length $L = 6$ m and a four-bladed rotor centered on the cylinder axis. The rotor cross-section has four-fold symmetry with a bounding diameter of 2 m, and its profile is twisted by 100° along the axial (z -)direction, producing a helical geometry representative of axial-flow machinery.

A uniform axial inflow $v_{in} = 20$ m/s is prescribed at $z = 0$, no-slip conditions are imposed on the cylindrical wall, and a stress-free condition is applied at the outlet $z = L$. On the rotor surface, the velocity is set by rigid-body rotation about the z -axis,

$$\mathbf{u}^f|_{\Gamma_{\text{rotor}}} = \dot{\theta} \mathbf{e}_z \times (\mathbf{x} - \mathbf{x}_c), \quad (50)$$




Fig. 27. Setup of the 3D rotating-rotor benchmark. (a) Cross-section through the cylindrical casing showing the four-bladed rotor profile and the annular fluid region. (b) Longitudinal section along the z -axis: the rotor blade envelope is twisted by 100° between inlet and outlet, and uniform axial inflow/stress-free outflow conditions are imposed at the two ends.

with angular velocity $\dot{\theta} = 30^\circ/\text{s}$. The kinematic viscosity is $\nu^f = 0.6 \text{ m}^2/\text{s}$, giving a Reynolds number $\text{Re} = v_{\text{in}} L / \nu^f = 200$. Backward Euler time integration is used with a constant step $\Delta t = 0.05 \text{ s}$ over the interval $[0, 5] \text{ s}$.

As in the two-dimensional cases, the fluid domain is decomposed into four ring-shaped patches, and the barrier-patch parameterization is reconstructed at each time step to update the volumetric mesh to the current rotor configuration.

Results. Fig. 28 shows representative snapshots at $t = 2.5 \text{ s}$ (after a 75° rotation) and $t = 5.0 \text{ s}$ (after 150°). At both instants, the four volumetric patches conform to the helical rotor surface (panels a, b), and the velocity field exhibits the expected acceleration through the blade-tip passages (panels c, d). Two complementary cross-sectional views are included: a vertical cut through the cylinder axis (panels e, f), which reveals the annular velocity distribution around the four-lobed rotor profile, and a constant- z cut near the inlet (panels g, h), where the four-fold symmetric pattern is seen to rotate with the rotor.

The tangential-slip reparameterization introduced in Section 3.2.3 is applied in the same way as in the two-dimensional rotating-square case, despite the substantially more demanding geometric setting: the four ring-shaped patches must accommodate a 100° helical twist along the axis while remaining conforming to the rotor surface as it rotates. Over the simulated interval $[0, 5] \text{ s}$, the rotor undergoes a cumulative rotation of 150° and the parameterization is updated at every one of the 100 time steps without a single failure of the barrier-patch parameterization. This indicates that the proposed parameterization-driven strategy generalizes to 3D rotating-machinery configurations with no algorithmic modifications beyond the natural extension of the patch topology.

6. Conclusions and outlook

We have presented a parameterization-driven approach to mesh motion in isogeometric ALE-FSI. Instead of deforming a reference mesh from one time step to the next, the fluid mesh is obtained at each step as a multi-patch spline parameterization built from the current interface geometry. The method relies on three ingredients: a barrier-based parameterization procedure that keeps the Jacobian positive; a tangential-slip reparameterization that allows the parametric pre-images to move along the interface, which is needed when the body rotates through large cumulative angles; and a constant-preserving quasi-interpolation operator that transfers the discrete solution between consecutive parameterizations while satisfying the discrete geometric conservation law at the algebraic level. The proposed method works for meshes defined by elements of any order. However, our analysis employs IGA-native spline-based mesh representations, which provide typically higher resistance to large mesh deformations, as we have demonstrated in Section 4.1.

On the perpendicular-flap and Turek–Hron FSI2 benchmarks, the method reproduces established reference results and keeps the Jacobian bounded away from zero for the full duration of the simulations. The rotating-square benchmark, in which classical mesh-deformation ALE fails after a limited rotation, is handled without difficulty by the combination of barrier-based parameterization and tangential slip. The prescribed falling-sphere test completes a 40 m translation with a positive Jacobian, zero reconnections, and zero mesh-size growth; the corresponding reference’s reconnection strategy increases the node and element counts by approximately 10%. The coupled free-fall test reproduces the wall-corrected terminal speed with a relative difference of 0.29%. The three-dimensional rotating-rotor example shows that the same construction applies to volumetric spline parameterizations. Since the parameterization is a purely geometric object, we also used the sequence of G+Smo-generated parameterizations to drive a deal.II finite element computation and obtained results consistent with the isogeometric solution.

The current framework has three limitations worth stating. The tangential-slip strategy, as implemented here, targets closed boundaries of convex or nearly convex bodies; strongly non-convex or multi-body configurations require a more general criterion for redistributing the pre-images. As demonstrated by the falling-sphere tests, large translational motion can be handled while the topology of the fluid domain and the patch connectivity remain unchanged. In contrast, topology-changing events such as collision, merging, or splitting of bodies are outside the scope of the present formulation, since they require a change of the patch layout itself.




Fig. 28. Three-dimensional rotating-rotor benchmark at $t = 2.5$ s (top row) and $t = 5.0$ s (bottom row); each column corresponds to one view type. Panels (a, e) show the four conforming volumetric patches of the parameterization adapted to the helically twisted rotor. Because the outer cylinder wall would otherwise occlude the interior flow, the velocity field is shown on three complementary views: panels (b, f) display the magnitude on the rotor surface together with the outer cylinder; panels (c, g) use a vertical cut aligned with the cylinder axis to reveal the velocity distribution around the four-lobed rotor profile; and panels (d, h) use a horizontal cut near the inlet that exposes the four-fold symmetric flow pattern on a constant- z plane. All velocity panels share the common color bar shown at the top.

Such cases would require repartitioning or a coupling with unfitted methods. The quasi-interpolation operator preserves constants but is not conservative for integral quantities such as mass or momentum, which matters for long-time simulations with strong mass transport and would call for a projection-based transfer instead.

A natural next step is to extend the tangential-slip idea to the non-convex and multi-body cases mentioned above, and to combine the per-step parameterization with adaptive or quality-driven strategies. More broadly, because the parameterization is built independently of the field solver, the same geometry module can drive other field solvers: the DEAL.II computation of Section 4.5 is a first demonstration of this, and we intend to extend it beyond FSI to other multi-physics problems with moving interfaces.

CRediT authorship contribution statement

Jingya Li: Writing – original draft, Visualization, Software, Methodology, Formal analysis, Investigation, Validation, Conceptualization. **Ye Ji:** Writing – original draft, Writing – review & editing, Methodology, Supervision, Visualization, Validation, Software, Investigation, Conceptualization. **Hugo Verhelst:** Writing – review & editing, Supervision, Software, Funding acquisition. **Henk den Besten:** Writing – review & editing, Supervision, Funding acquisition. **Matthias Möller:** Writing – review & editing, Supervision, Software, Project administration, Funding acquisition, Conceptualization.

Funding

This research has been conducted as part of the project FlexFloat with project number 19002 of the research program Open Technology, which is (partly) financed by the Dutch Research Council (NWO), Netherlands.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

The authors gratefully acknowledge Hana Horníková and Bohumír Bastl from the University of West Bohemia in Pilsen, Czech Republic, for generously sharing their IGA fluid solver and for their helpful supports throughout this work. We also thank Teresa Schalk and Thomas-Peter Fries at Graz University of Technology for providing the rotating-square benchmark parameters and reference kinetic-energy and momentum data used in the validation.

Data availability

Data will be made available on request.

References

- [1] N.D. Katopodes, *Free-Surface Flow: Computational Methods*, Butterworth-Heinemann, 2018.
- [2] T.E. Tezduyar, Finite element methods for flow problems with moving boundaries and interfaces, *Arch. Comput. Methods Eng.* 8 (2) (2001) 83–130, <http://dx.doi.org/10.1007/BF02897870>.
- [3] Y. Bazilevs, K. Takizawa, T.E. Tezduyar, *Computational Fluid–Structure Interaction: Methods and Applications*, in: Wiley Series in Computational Mechanics, John Wiley & Sons, 2013, <http://dx.doi.org/10.1002/9781118483565>.
- [4] G. Hou, J. Wang, A. Layton, Numerical methods for fluid-structure interaction – a review, *Commun. Comput. Phys.* 12 (2) (2012) 337–377, <http://dx.doi.org/10.4208/cicp.291210.290411s>.
- [5] S. Takagi, K. Sugiyama, S. Ii, Y. Matsumoto, A review of full Eulerian methods for fluid structure interaction problems, *J. Appl. Mech.* 79 (1) (2012) 010911, <http://dx.doi.org/10.1115/1.4005184>.
- [6] C.S. Peskin, The immersed boundary method, *Acta Numer.* 11 (2002) 479–517, <http://dx.doi.org/10.1017/S0962492902000077>.
- [7] B. Schott, C. Ager, W.A. Wall, Monolithic cut finite element-based approaches for fluid-structure interaction, *Internat. J. Numer. Methods Engrg.* 119 (8) (2019) 757–796, <http://dx.doi.org/10.1002/nme.6072>.
- [8] E. Burman, P. Hansbo, M.G. Larson, S. Zahedi, Cut finite element methods, *Acta Numer.* 34 (2025) 1–121, <http://dx.doi.org/10.1017/S0962492925000017>.
- [9] B. Marussig, T.J.R. Hughes, A review of trimming in isogeometric analysis: Challenges, data exchange and simulation aspects, *Arch. Comput. Methods Eng.* 25 (4) (2018) 1059–1127, <http://dx.doi.org/10.1007/s11831-017-9220-9>.
- [10] D. Kamensky, M.-C. Hsu, D. Schillinger, J.A. Evans, A. Aggarwal, Y. Bazilevs, M.S. Sacks, T.J.R. Hughes, An immersogeometric variational framework for fluid-structure interaction: application to bioprosthetic heart valves, *Comput. Methods Appl. Mech. Engrg.* 284 (2015) 1005–1053, <http://dx.doi.org/10.1016/j.cma.2014.10.040>.
- [11] N. Hosters, J. Helmig, A. Stavrev, M. Behr, S. Elgeti, Fluid-structure interaction with NURBS-based coupling, *Comput. Methods Appl. Mech. Engrg.* 332 (2018) 520–539, <http://dx.doi.org/10.1016/j.cma.2018.01.003>.
- [12] J.-Y. Li, H.M. Verhelst, H. den Besten, M. Möller, Isogeometric suitable coupling methods for partitioned multiphysics simulation with application to fluid-structure interaction, *Eng. Comput.* 42 (2) (2026) 69.
- [13] C.W. Hirt, A.A. Amsden, J.L. Cook, An arbitrary Lagrangian-Eulerian computing method for all flow speeds, *J. Comput. Phys.* 14 (3) (1974) 227–253, [http://dx.doi.org/10.1016/0021-9991\(74\)90051-5](http://dx.doi.org/10.1016/0021-9991(74)90051-5).
- [14] J. Donea, A. Huerta, J.-P. Ponthot, A. Rodriguez-Ferran, Arbitrary Lagrangian-Eulerian methods, in: *Encyclopedia of Computational Mechanics*, John Wiley & Sons, 2004, pp. 413–437, <http://dx.doi.org/10.1002/0470091355.ecm009>.
- [15] A.A. Johnson, T.E. Tezduyar, Mesh update strategies in parallel finite element computations of flow problems with moving boundaries and interfaces, *Comput. Methods Appl. Mech. Engrg.* 119 (1–2) (1994) 73–94, [http://dx.doi.org/10.1016/0045-7825\(94\)00077-8](http://dx.doi.org/10.1016/0045-7825(94)00077-8).
- [16] B.T. Helenbrook, Mesh deformation using the biharmonic operator, *Internat. J. Numer. Methods Engrg.* 56 (7) (2003) 1007–1021, <http://dx.doi.org/10.1002/nme.595>.
- [17] T.E. Tezduyar, M. Behr, J. Liou, A new strategy for finite element computations involving moving boundaries and interfaces—the deforming-spatial-domain/space-time procedure: I. The concept and the preliminary numerical tests, *Comput. Methods Appl. Mech. Engrg.* 94 (3) (1992) 339–351, [http://dx.doi.org/10.1016/0045-7825\(92\)90059-S](http://dx.doi.org/10.1016/0045-7825(92)90059-S).
- [18] F.J. Blom, Considerations on the spring analogy, *Internat. J. Numer. Methods Fluids* 32 (6) (2000) 647–668, [http://dx.doi.org/10.1002/\(SICI\)1097-0363\(20000330\)32:6<647::AID-FLD979>3.0.CO;2-K](http://dx.doi.org/10.1002/(SICI)1097-0363(20000330)32:6<647::AID-FLD979>3.0.CO;2-K).
- [19] C. Farhat, P. Geuzaine, C. Grandmont, The discrete geometric conservation law and the nonlinear stability of ALE schemes for the solution of flow problems on moving grids, *J. Comput. Phys.* 174 (2) (2001) 669–694, <http://dx.doi.org/10.1006/jcph.2001.6932>.
- [20] L. Formaggia, F. Nobile, Stability analysis of second-order time accurate schemes for ALE–FEM, *Comput. Methods Appl. Mech. Engrg.* 193 (39–41) (2004) 4097–4116, <http://dx.doi.org/10.1016/j.cma.2003.09.028>.
- [21] T. Hughes, J. Cottrell, Y. Bazilevs, Isogeometric analysis: CAD, finite elements, NURBS, exact geometry and mesh refinement, *Comput. Methods Appl. Mech. Engrg.* 194 (39) (2005) 4135–4195, <http://dx.doi.org/10.1016/j.cma.2004.10.008>, URL <https://www.sciencedirect.com/science/article/pii/S0045782504005171>.
- [22] Y. Bazilevs, V.M. Calo, T.J.R. Hughes, Y. Zhang, Isogeometric fluid-structure interaction: theory, algorithms, and computations, *Comput. Mech.* 43 (1) (2008) 3–37, <http://dx.doi.org/10.1007/s00466-008-0315-x>.
- [23] Y. Bazilevs, V.M. Calo, Y. Zhang, T.J.R. Hughes, Isogeometric fluid–structure interaction analysis with applications to arterial blood flow, *Comput. Mech.* 38 (4) (2006) 310–322, <http://dx.doi.org/10.1007/s00466-006-0084-3>.
- [24] Y. Bazilevs, K. Takizawa, T.E. Tezduyar, ALE-VMS and ST-VMS methods for computer modeling of wind-turbine rotor aerodynamics and fluid-structure interaction, *Math. Models Methods Appl. Sci.* 23 (2) (2013) 249–267, <http://dx.doi.org/10.1142/S0218202513400034>.
- [25] Y. Bazilevs, K. Takizawa, T.E. Tezduyar, A. Korobenko, T. Kuraishi, Y. Otoguro, Computational aerodynamics with isogeometric analysis, *J. Mech.* 39 (2023) 24–39, <http://dx.doi.org/10.1093/jom/ufad002>.
- [26] P. Tonon, R.A.K. Sanches, K. Takizawa, T.E. Tezduyar, A linear-elasticity-based mesh moving method with no cycle-to-cycle accumulated distortion, *Comput. Mech.* 67 (2) (2021) 413–434, <http://dx.doi.org/10.1007/s00466-020-01941-y>.

- [27] R. Codina, G. Houzeaux, H. Coppola-Owen, J. Baiges, The fixed-mesh ALE approach for the numerical approximation of flows in moving domains, *J. Comput. Phys.* 228 (5) (2009) 1591–1611, <http://dx.doi.org/10.1016/j.jcp.2008.11.004>.
- [28] J. Baiges, R. Codina, The fixed-mesh ALE approach applied to solid mechanics and fluid-structure interaction problems, *Internat. J. Numer. Methods Engrg.* 81 (12) (2010) 1529–1557, <http://dx.doi.org/10.1002/nme.2740>.
- [29] T. Schwentner, T.-P. Fries, Fully coupled, higher-order, block-structured mesh generation in fluid-structure interaction, *Internat. J. Numer. Methods Fluids* 97 (3) (2025) 359–377, <http://dx.doi.org/10.1002/fld.5355>.
- [30] T. Schalk, T.-P. Fries, Block-structured meshing in coupled problems with moving domains, *Proc. Appl. Math. Mech.* 26 (2) (2026) e70110, <http://dx.doi.org/10.1002/pamm.70110>.
- [31] D. Di Cristofaro, A. Frangi, M. Cremonesi, 3D fluid–structure interaction simulation with an Arbitrary-Lagrangian–Eulerian approach with applications to flying objects, *Eng. Comput.* 41 (3) (2025) 1747–1769, <http://dx.doi.org/10.1007/s00366-024-02043-6>.
- [32] P.E. Farrell, J.R. Maddison, Conservative interpolation between volume meshes by local Galerkin projection, *Comput. Methods Appl. Mech. Engrg.* 200 (1–4) (2011) 89–100, <http://dx.doi.org/10.1016/j.cma.2010.07.015>.
- [33] A. Pont, R. Codina, J. Baiges, Interpolation with restrictions between finite element meshes for flow problems in an ALE setting, *Internat. J. Numer. Methods Engrg.* 110 (13) (2017) 1203–1226, <http://dx.doi.org/10.1002/nme.5444>.
- [34] Y. Ji, Y.-Y. Yu, M.-Y. Wang, C.-G. Zhu, Constructing high-quality planar NURBS parameterization for isogeometric analysis by adjustment control points and weights, *J. Comput. Appl. Math.* 396 (2021) 113615, <http://dx.doi.org/10.1016/j.cam.2021.113615>.
- [35] L. Piegl, W. Tiller, *The NURBS Book*, Springer Science & Business Media, 2012.
- [36] B. Sayyid Hosseini, M. Möller, S. Turek, Isogeometric Analysis of the Navier–Stokes equations with Taylor–Hood B-spline elements, *Appl. Math. Comput.* 267 (2014) 264–281, <http://dx.doi.org/10.1016/j.amc.2015.03.104>.
- [37] Y. Ji, M.-Y. Wang, M.-D. Pan, Y. Zhang, C.-G. Zhu, Penalty function-based volumetric parameterization method for isogeometric analysis, *Comput. Aided Geom. Design* 94 (2022) 102081, <http://dx.doi.org/10.1016/j.cagd.2022.102081>.
- [38] Y. Ji, M. Möller, Y. Yu, C. Zhu, Boundary parameter matching for isogeometric analysis using Schwarz–Christoffel mapping, *Eng. Comput.* 40 (6) (2024) 3929–3947.
- [39] A. Buffa, E.M. Garau, C. Giannelli, G. Sangalli, On quasi-interpolation operators in spline spaces, in: *Building Bridges: Connections and Challenges in Modern Approaches to Numerical Partial Differential Equations*, Springer, 2016, pp. 73–91.
- [40] H. Guillard, C. Farhat, On the significance of the geometric conservation law for flow computations on moving meshes, *Comput. Methods Appl. Mech. Engrg.* 190 (11–12) (2000) 1467–1482, [http://dx.doi.org/10.1016/S0045-7825\(00\)00173-0](http://dx.doi.org/10.1016/S0045-7825(00)00173-0).
- [41] H. Casquero, C. Bona-Casas, H. Gomez, A NURBS-based immersed methodology for fluid–structure interaction, *Comput. Methods Appl. Mech. Engrg.* 284 (2015) 943–970, <http://dx.doi.org/10.1016/j.cma.2014.10.055>.
- [42] R.T. Black, S.A. Maas, W. Wu, J. Maheshwari, T. Kolev, J.A. Weiss, M.A. Jolley, An open-source computational framework for immersed fluid–structure interaction modeling using FEBio and MFEM, *Eng. Comput.* 42 (2026) 166, <http://dx.doi.org/10.1007/s00366-026-02371-9>.
- [43] B. Jüttler, U. Langer, A. Mantzaflaris, S.E. Moore, W. Zulehner, Geometry + simulation modules: Implementing isogeometric analysis, *Proc. Appl. Math. Mech.* 14 (1) (2014) 961–962, <http://dx.doi.org/10.1002/pamm.201410461>.
- [44] A. Mantzaflaris, H.M. Verhelst, M. Möller, C. Karampatzakis, S. Takacs, S. Imperatore, P. Weinmüller, J. Vogl, J. Li, Y. Ji, M. Matucci, A. Limkilde, F. Scholz, D. Mokriš, J. Sogn, A. Shamanskiy, E. Vollebregt, R. Tielen, A. Bressan, A. Farahat, *Geometry + Simulation Modules (G+Smo)*, 2025, <http://dx.doi.org/10.5281/zenodo.15874172>, Version v25.07.0.
- [45] D. Arndt, W. Bangerth, M. Bergbauer, M. Feder, M. Fehling, J. Heinz, T. Heister, L. Heltai, M. Kronbichler, M. Maier, P. Munch, J.-P. Pelteret, B. Turcksin, D. Wells, S. Zampini, The deal.II library, version 9.5, *J. Numer. Math.* 31 (3) (2023) 231–246, <http://dx.doi.org/10.1515/jnma-2023-0089>.
- [46] G. Chourdakis, K. Davis, B. Rodenberg, M. Schulte, F. Simonis, B. Uekermann, G. Abrams, H.-J. Bungartz, L.C. Yau, I. Desai, K. Eder, R. Hertrich, F. Lindner, A. Rusch, D. Sashko, D. Schneider, A. Totounferoush, D. Volland, P. Vollmer, O.Z. Koseomur, preCICE v2: A sustainable and user-friendly coupling library, *Open Res. Eur.* 2 (2022) 51, <http://dx.doi.org/10.12688/openreseurope.14445.2>.
- [47] T. Wick, Fluid-structure interactions using different mesh motion techniques, *Comput. Struct.* 89 (13–14) (2011) 1456–1467, <http://dx.doi.org/10.1016/j.compstruc.2011.02.019>.
- [48] K. Stein, T. Tezduyar, Advanced mesh update techniques for problems involving large displacements, in: *Proceedings of the Fifth World Congress on Computational Mechanics*, Vienna, Austria, 2002, pp. Paper-ID 81489, URL <http://wccm.tuwien.ac.at/>.
- [49] A. Shamanskiy, B. Simeon, Mesh moving techniques in fluid-structure interaction: robustness, accumulated distortion and computational efficiency, *Comput. Mech.* 67 (2) (2021) 583–600, <http://dx.doi.org/10.1007/s00466-020-01950-x>.
- [50] N. Delaissé, T. Demeester, R. Haelterman, J. Degroote, Quasi-Newton methods for partitioned simulation of fluid-structure interaction reviewed in the generalized Broyden framework, *Arch. Comput. Methods Eng.* 30 (5) (2023) 3271–3300, <http://dx.doi.org/10.1007/s11831-023-09907-y>.
- [51] S. Turek, J. Hron, Proposal for numerical benchmarking of fluid-structure interaction between an elastic object and laminar incompressible flow, in: *Fluid-Structure Interaction: Modelling, Simulation, Optimisation*, Springer, 2006, pp. 371–385.
- [52] U. Küttler, W. Wall, Fixed-point fluid-structure interaction solvers with dynamic relaxation, *Comput. Mech.* 43 (1) (2008) 61–72, <http://dx.doi.org/10.1007/s00466-008-0255-5>.
- [53] J. Egermaier, H. Horníková, Approximate inner solvers for block preconditioning of the incompressible Navier–Stokes problems discretized by isogeometric analysis, *Internat. J. Numer. Methods Fluids* 96 (6) (2024) 1078–1103, <http://dx.doi.org/10.1002/fld.5280>.