

A comparison of non-matching techniques for the finite element approximation of interface problems

Daniele Boffi ^{a,b,1}, Andrea Cangiani ^{c,1}, Marco Feder ^{c,2}, Lucia Gastaldi ^{d,e,1,*}, Luca Heltai ^{c,1}

^a CEMSE division, King Abdullah University of Science and Technology (KAUST), Thuwal, Saudi Arabia

^b Dipartimento di Matematica 'F. Casorati', Università degli Studi di Pavia, Pavia, Italy

^c Mathematics Area, SISSA, Trieste, Italy

^d DICATAM Sez. di Matematica, Università degli Studi di Brescia, Brescia, Italy

^e Istituto di Matematica Applicata e Tecnologie Informatiche (IMATI-CNR), Pavia, Italy

ARTICLE INFO

Dedicated to Leszek Demkowicz on the occasion of his 70th birthday.

Keywords:

Interface problems
Finite element method
Non-matching techniques
Lagrange multiplier
Nitsche method

ABSTRACT

We perform a systematic comparison of various numerical schemes for the approximation of interface problems. We consider unfitted approaches in view of their application to possibly moving configurations. Particular attention is paid to the implementation aspects and to the analysis of the costs related to the different phases of the simulations.

1. Introduction

The efficient numerical solution of partial differential equations modeling the interaction of physical phenomena across interfaces with complex, possibly moving, shapes is of great importance in many scientific fields. We refer, for instance, to fluid-structure interaction, or crack propagation, just to mention two relevant examples. A crucial issue is the handling of computational grids. In this respect, we can classify computational methods for interface problems into two families: boundary fitted methods and boundary unfitted methods. For time dependent problems, the former are typically handled using the Arbitrary Eulerian Lagrangian formulation ([13], [22]), where meshes are deformed in a conforming way with respect to movements of the physical domains. In this case, the imposition of interface conditions is usually easy to implement. However, an accurate description of both meshes is required, and the allowed movements are restricted by the topological structure of the initial state. When topology may change, or when the grid undergoes severe deformations, these methods require remeshing. An operation which is computationally very expensive in time dependent scenarios and three dimensional settings. Conversely, unfitted

approaches are based on describing the physical domains as embedded into a constant background mesh. As it does not require remeshing, this approach is extremely flexible, but it requires sophisticated methods to represent interfaces. Among unfitted approaches, we mention the Immersed Boundary Method [7,28], the Cut Finite Element Method [11], and the Extended Finite Element Method [26]. Another popular choice is the Fictitious Domain Method with Lagrange multipliers, proposed by Glowinski, Pan and Periaux in [17] for a Dirichlet problem, analyzed in [15], and then extended to particulate flows in [16].

The use of a Lagrange multiplier for dealing with Dirichlet boundary conditions was introduced in the seminal work by Babuška [3]. This is formulated in terms of a symmetric saddle point problem where the condition at the interface is enforced through the use of a Lagrange multiplier. The main drawback of this method is that it suffers from a loss of accuracy at interfaces, even if it is known [20] that this detrimental effect on the convergence properties of the approximate solution is a local phenomenon, restricted to a small neighborhood of the interface. From the computational standpoint, the Fictitious Domain Method poses the additional challenge of assembling coupling terms involving basis functions living on different meshes. In this context, one can dis-

* Corresponding author.

E-mail address: lucia.gastaldi@unibs.it (L. Gastaldi).

¹ Contributing author, provided significant input and contributions to the manuscript.

² Lead author, responsible for the overall direction and content of the manuscript.

tribute quadrature points on the immersed mesh and let it drive the integration process, or one may compute a composite quadrature rule by identifying exact intersections (polytopes, in general) between the two meshes. This approach has been presented in different papers and frameworks: for instance, in [24], a high performance library has been developed to perform such tasks, and in [5] it has been shown how composite rules on interfaces turn out to be necessary to recover optimal rates for a fluid-structure interaction problem where both the solid and the fluid meshes are two dimensional objects. In the cut-FEM framework, we mention [25] for an efficient implementation of Nitsche’s method on three dimensional overlapping meshes.

In this paper we consider a two or three dimensional domain with an immersed interface of co-dimension one and we study the numerical approximation of the solution of an elliptic PDE whose solution is prescribed along the interface. Although in our case the domain and the interface are fixed, we are discussing the presented numerical schemes in view of their application to more general settings. We perform a systematic comparison between three different unfitted approaches, analyzing them in terms of accuracy, computational cost, and implementation effort. In particular, we perform a comparative analysis in terms of accuracy and CPU times for the Lagrange multiplier method, Nitsche’s penalization method, and cut-FEM for different test cases, and discuss the benefit of computing accurate quadrature rules on mesh intersections.

The outline of the paper is as follows. In Section 2 we introduce the model Poisson problem posed on a domain with an internal boundary and in Section 3 we review the Lagrange multiplier method, the Nitsche penalization method, and the cut-FEM method for its solution. In Section 4 we discuss the central issue of the numerical integration of the coupling terms, while in Section 5 we present a numerical comparison of the three methods, focusing on the validation of the implementation, on how coupling terms affect the accuracy of numerical solutions, and on computational times. Finally, the results are summarised in Section 6.

2. Model problem and notation

Let ω be a closed and bounded domain of \mathbb{R}^d , $d = 2, 3$, with Lipschitz continuous boundary $\gamma := \partial\omega$, and $\Omega \subset \mathbb{R}^d$ a Lipschitz domain such that $\omega \Subset \Omega$; see Fig. 1 for a prototypical configuration. We consider the model problem

$$\begin{cases} -\Delta u = f & \text{in } \Omega \setminus \gamma, \\ u = g & \text{on } \gamma, \\ u = 0 & \text{on } \Gamma := \partial\Omega, \end{cases} \quad (1)$$

for given data $f \in L^2(\Omega)$ and $g \in H^{\frac{1}{2}}(\gamma)$. Throughout this work we refer to Ω as the *background* domain, while we refer to ω as the *immersed* domain, and γ as the *immersed boundary*. The rationale behind this setting is that it allows to solve problems in a complex and possibly time dependent domain ω , by embedding the problem in a simpler background domain Ω – typically a box – and imposing some constraints on the immersed boundary γ . For the sake of simplicity, we consider the case in which the immersed domain is *entirely* contained in the background domain, but more general configurations may be considered.

As ambient spaces for (1), we consider $V(\Omega) := H_0^1(\Omega) = \{v \in H^1(\Omega) : v|_{\Gamma} = 0\}$ and $Q(\gamma) := H^{-\frac{1}{2}}(\gamma)$. Given a domain $D \subset \mathbb{R}^d$ and a real number $s \geq 0$, we denote by $\|\cdot\|_{s,D}$ the standard norm of $H^s(D)$. In particular, $\|\cdot\|_{0,D}$ stands for the L^2 -norm stemming from the standard L^2 -inner product $(\cdot, \cdot)_D$ on D . Finally, with $\langle \cdot, \cdot \rangle_\gamma$ we denote the standard duality pairing between $Q(\gamma)$ and its dual $Q'(\gamma) = H^{\frac{1}{2}}(\gamma)$.

Across the immersed boundary γ , we define the *jump* operator as

$$\begin{aligned} \llbracket v \rrbracket_\gamma &= v^+ - v^-, \\ \llbracket \tau \rrbracket_\gamma &= \tau^+ - \tau^-, \end{aligned}$$

for smooth enough scalar- and vector- valued functions v and τ . Here, v^\pm and τ^\pm are external and internal traces defined according to the direction of the outward normal \mathbf{n} to ω at γ .

Problem (1) can be written as a constrained minimization problem by introducing the Lagrangian $\mathcal{L} : V(\Omega) \times Q(\gamma) \rightarrow \mathbb{R}$ defined as

$$\mathcal{L}(v, q) := \frac{1}{2}(\nabla v, \nabla v)_\Omega - (f, v)_\Omega + \langle q, v - g \rangle_\gamma. \quad (2)$$

Looking for stationary points of \mathcal{L} gives the following saddle point problem of finding a pair $(u, \lambda) \in V(\Omega) \times Q(\gamma)$ such that

$$\begin{aligned} (\nabla u, \nabla v)_\Omega + \langle \lambda, v \rangle_\gamma &= (f, v)_\Omega \quad \forall v \in V(\Omega), \\ \langle q, u \rangle_\gamma &= \langle q, g \rangle_\gamma \quad \forall q \in Q(\gamma). \end{aligned} \quad (3)$$

Below in Theorem 1 we show that this problem admits a unique solution. Starting from (1) and integrating by parts, one can easily show that setting $\lambda = -\llbracket \nabla u \cdot \mathbf{n} \rrbracket_\gamma$, the pair $(u, \lambda) \in V(\Omega) \times Q(\gamma)$ is the solution of (3)-(4). Conversely, with proper choices for $v \in V(\Omega)$ in (3) one gets that $-\Delta u = f$ in $\Omega \setminus \gamma$, while (4) implies $u = g$ on γ , so that (3)-(4) are equivalent to (1) with $-\lambda$ equal to the jump of the normal derivative of u on the interface γ . Moreover, if the datum g is sufficiently smooth, say $g \in H^s(\gamma)$ for $s > 1$, we can further take $\lambda \in L^2(\gamma)$ and use in practice $Q(\gamma) = L^2(\gamma)$.

In the following theorem we sketch the proof of existence and uniqueness of the solution of (3)-(4).

Theorem 1. *Given $f \in L^2(\Omega)$ and $g \in H^{1/2}(\gamma)$, there exists a unique solution of Problem (3)-(4) satisfying the following stability estimate*

$$\|u\|_{1,\Omega} + \|\lambda\|_{-1/2,\gamma} \leq C(\|f\|_{0,\Omega} + \|g\|_{1/2,\gamma}).$$

Proof. Problem (3)-(4) is a saddle point problem, hence to show existence and uniqueness of the solution we need to check the *ellipticity on the kernel* and the *inf-sup condition* [4]. The kernel $\mathbb{K} = \{v \in V(\Omega) : \langle q, v \rangle = 0 \ \forall q \in Q(\gamma)\}$, can be identified with the subset of functions in $V(\Omega)$ with vanishing trace along γ . Then thanks to the Poincaré inequality we have that there exists $\alpha_0 > 0$ such that

$$(\nabla u, \nabla u)_\Omega \geq \alpha_0 \|u\|_{1,\Omega}^2.$$

The inf-sup condition can be verified using the definition of the norm in $Q(\gamma)$ and the fact that, by the trace theorem, for each $w \in H^{1/2}(\gamma)$ there exists at least an element $v \in V(\Omega)$ such that $v = w$ on γ , with $\|v\|_{1,\Omega} \leq C_1 \|w\|_{1/2,\gamma}$. Hence we get the inf-sup condition

$$\inf_{q \in Q(\gamma)} \sup_{v \in V(\Omega)} \frac{\langle q, v \rangle}{\|q\|_{-1/2,\gamma} \|v\|_{1,\Omega}} \geq \beta_0,$$

with $\beta_0 = 1/C_1$. \square

A detailed analysis for the Dirichlet problem, where the Lagrange multiplier is used to impose the boundary condition, can be found in the pioneering work by Glowinski, Pan and Periaux [17].

3. Non-matching discretizations

We assume that both Ω and ω are Lipschitz domains and we discretize the problem introducing computational meshes for the domain Ω and for the immersed boundary γ which are *unfitted* with respect to each other in that they are constructed independently. The computational meshes Ω_h of Ω and γ_h of γ consist of disjoint elements such that $\Omega = \bigcup_{T \in \Omega_h} T$ and $\gamma = \bigcup_{K \in \gamma_h} K$. When $d = 2$, Ω_h will be a triangular or quadrilateral mesh and γ_h a mesh composed by straight line segments. For $d = 3$, Ω_h will be a tetrahedral or hexahedral mesh and γ_h a surface mesh whose elements are triangles or planar quadrilaterals embedded in the three dimensional space. We denote by h_Ω and h_γ the mesh sizes of Ω_h and γ_h , respectively. For simplicity, we ignore geometrical errors in the discretizations of Ω , and γ , and we assume that the mesh sizes

h_Ω and h_γ are small enough so that the geometrical error is negligible with respect to the discretization error (see Sect. 4.3 for a quantitative estimate of the geometrical error in our numerical experiments).

We consider discretizations of $V(\Omega)$ based on standard Lagrange finite elements, namely

$$V_h(\Omega) := \{v \in H_0^1(\Omega) : v|_T \in \mathcal{R}^p(T), \forall T \in \mathcal{T}_h\} \quad p \geq 1, \tag{5}$$

with $\mathcal{R}^p(T) := \mathcal{P}^p(T)$ or $\mathcal{Q}^p(T)$, the spaces of polynomials of total degree up to p or of degree p separately in each variable, respectively, depending on whether T is a simplex or a quadrilateral/hexahedron. For the cut-FEM method, the corresponding spaces will be constructed separately in each subdomain; this results in a doubling of the degrees of freedom on the elements cut by the interface (see Sect. 3.3).

For the Lagrange multiplier space $Q(\gamma)$ we employ the space of piecewise-polynomial functions

$$Q_h(\gamma) := \{q \in L^2(\gamma) : q|_K \in \mathcal{R}^p(K), \forall K \in \mathcal{T}_h\} \quad p \geq 0, \tag{6}$$

which we equip with the mesh dependent norm

$$\|q\|_{-\frac{1}{2},\gamma} := \|h^{-\frac{1}{2}}q\|_{0,\gamma} \quad \forall q \in Q_h(\gamma), \tag{7}$$

where h is the piecewise constant function given by $h|_K = h_K$ the diameter of K for each $K \in \mathcal{T}_h$. The definition and the notation are justified by the fact that, on a quasi-uniform mesh, the mesh dependent norm is equivalent to the norm in $H^{1/2}(\gamma)$.

3.1. The method of Lagrange multipliers

The discrete counterpart of (3)-(4) is to find a pair $(u_h, \lambda_h) \in V_h \times Q_h$ such that

$$(\nabla u_h, \nabla v_h)_\Omega + \langle \lambda_h, v_h \rangle_\gamma = (f, v_h)_\Omega \quad \forall v_h \in V_h, \tag{8}$$

$$\langle q_h, u_h \rangle_\gamma = \langle q_h, g \rangle_\gamma \quad \forall q_h \in Q_h. \tag{9}$$

The next theorem states existence, uniqueness, and stability of the discrete solution together with optimal error estimates.

Theorem 2. Assume that the mesh γ_h is quasi uniform and that there exists a constant $C_r > 1$ independent of h_Ω and h_γ such that $h_\Omega/h_\gamma \leq C_r$. Then, there exists a unique solution $(u_h, \lambda_h) \in V_h \times Q_h$ of Problem (8)-(9). Moreover, it holds

$$\|u - u_h\|_{1,\Omega} + \|\lambda - \lambda_h\|_{-1/2,\gamma} \leq C \inf_{\substack{v_h \in V_h \\ \mu_h \in Q_h}} (\|u - v_h\|_{1,\Omega} + \|\lambda - \mu_h\|_{-1/2,\gamma}), \tag{10}$$

with $C > 0$ a constant independent of the mesh sizes h_Ω and h_γ .

Proof. The existence, uniqueness, and stability of the discrete solution can be obtained by showing that there exist positive constants α and β , independent of h_Ω and h_γ , such that the ellipticity on the kernel and inf-sup condition hold true at the discrete level [4]. Since $Q_h \subset L^2(\gamma)$, we have that the discrete kernel $\mathcal{K}_h = \{v_h \in V_h : \langle q_h, v_h \rangle = 0 \forall q_h \in Q_h\}$ contains element with $\int_\gamma v_h ds = 0$. Hence, the Poincaré inequality (see [9, (5.3.3)])

$$\|v_h\|_{1,\Omega} \leq C_\Omega \left(\left| \int_\gamma v_h ds \right| + \|\nabla v_h\|_{0,\Omega} \right) = C_\Omega \|\nabla v_h\|_{0,\Omega},$$

implies that

$$(\nabla v_h, \nabla v_h)_\Omega \geq \alpha \|v_h\|_{1,\Omega}^2 \quad \forall v_h \in \mathcal{K}_h.$$

The discrete inf-sup condition

$$\inf_{q_h \in Q_h} \sup_{v_h \in V_h} \frac{\langle q_h, v_h \rangle}{\|q_h\|_{-1/2,\gamma} \|v_h\|_{1,\Omega}} \geq \beta,$$




Fig. 1. Model problem setting, with immersed domain ω , immersed boundary γ , and background domain Ω .

is more involved and makes use of the continuous inf-sup, together with Clément interpolation, trace theorem and inverse inequality. The interested reader can find the main arguments of this proof in [6, sect. 5].

Thanks to the above conditions, the theory on the approximation of saddle point problems gives both existence and uniqueness of the solution of Problem (8)-(9) satisfying the a priori estimate

$$\|u_h\|_{1,\Omega} + \|\lambda_h\|_{-1/2,\gamma} \leq C(\|f\|_{0,\Omega} + \|g\|_{1/2,\gamma}),$$

and the error estimate (10). \square

Given basis functions $\{v_i\}_{i=1}^N$ and $\{q_i\}_{i=1}^M$ such that $V_h := \text{span}\{v_i\}_{i=1}^N$ and $Q_h := \text{span}\{q_i\}_{i=1}^M$, we have that (8), (9) can be written as the following algebraic problem

$$\begin{pmatrix} A & C^\top \\ C & 0 \end{pmatrix} \begin{pmatrix} U \\ \lambda \end{pmatrix} = \begin{pmatrix} F \\ G \end{pmatrix} \tag{11}$$

where

$$A_{ij} = (\nabla v_j, \nabla v_i)_\Omega \quad i, j = 1, \dots, N$$

$$C_{\alpha j} = \langle q_\alpha, v_j \rangle_\gamma \quad j = 1, \dots, N, \alpha = 1, \dots, M$$

$$F_i = (f, v_i)_\Omega \quad i = 1, \dots, N$$

$$G_\alpha = \langle q_\alpha, g \rangle_\gamma \quad \alpha = 1, \dots, M.$$

To solve the block linear system (11) we use Krylov subspace iterative methods applied to the Schur complement system:

$$\lambda = S^{-1}(CA^{-1}F - G), \tag{12}$$

$$U = A^{-1}(F - C^\top \lambda), \tag{13}$$

where $S := CA^{-1}C^\top$, and we use $CAC^\top + M$ as preconditioner for S , where M is the immersed boundary mass matrix with entries $(M)_{ij} = \langle q_j, q_i \rangle_\gamma$.

We next show how the Lagrange multiplier formulation is directly linked to a penalization approach used to impose the Dirichlet condition $u = g$ on the internal curve γ , by locally eliminating the multiplier in the same spirit of the work by Stenberg [30].

3.2. The method of Nitsche

Instead of enforcing the constraint on γ with a multiplier, it is possible to impose it weakly through a penalization approach following the so-called method of Nitsche. Here we show that the enforcement of boundary conditions via Nitsche's method can be derived from a stabilized Lagrange multiplier method by adding a consistent term that penalizes the distance between the discrete multiplier λ_h and the normal derivative [30]. With this in mind, we penalize the jump of the normal derivative along the internal curve γ to impose the constraint $u = g$. The consistency here follows from the observation that at the continuous level the multiplier is the jump of the normal derivative on the interface.

We define $h(x)$ as the piecewise constant function describing the mesh-size of γ and we choose the discrete spaces V_h and Q_h as in the Lagrange multiplier case.

Adding the normal gradient penalization term to the Lagrange multiplier formulation (8)-(9) leads to the problem of seeking a pair $(u_h, \lambda_h) \in V_h \times Q_h$ such that

$$\begin{aligned} (\nabla u_h, \nabla v_h)_\Omega + \langle \lambda_h, v_h \rangle_\gamma + \frac{1}{\beta} \langle \llbracket \nabla v_h \cdot \mathbf{n} \rrbracket, h(\lambda_h + \llbracket \nabla u_h \cdot \mathbf{n} \rrbracket) \rangle_\gamma &= (f, v_h)_\Omega \\ \forall v_h \in V_h, \\ \langle q_h, u_h \rangle_\gamma - \frac{1}{\beta} \langle q_h, h(\lambda_h + \llbracket \nabla u_h \cdot \mathbf{n} \rrbracket) \rangle_\gamma &= \langle q_h, g \rangle_\gamma \quad \forall q_h \in Q_h, \end{aligned}$$

where β is a positive penalty parameter. In the discrete setting, $\langle q, v \rangle_\gamma$ is identified with the scalar product in $L^2(\gamma)$ for $q \in Q_h$ and $v \in V_h$ and we use the notation $\langle q, v \rangle_\gamma = \sum_{K \in \gamma_h} \langle q, v \rangle_K$. The second equation gives

$$\langle q_h, u_h - \frac{1}{\beta} h(\lambda_h + \llbracket \nabla u_h \cdot \mathbf{n} \rrbracket) - g \rangle_\gamma = 0 \quad \forall q_h \in Q_h,$$

and, introducing the L^2 -projection on Q_h as $\Pi_h : L^2(\gamma) \rightarrow Q_h$, we can eliminate the multiplier locally on each element $K \in \gamma_h$:

$$\lambda_h|_K = -(\Pi_h \llbracket \nabla u_h \cdot \mathbf{n} \rrbracket)_K + \beta h_K^{-1} (\Pi_h (u_h - g))_K.$$

We observe that it is possible to formally refine γ_h to γ'_h such that the element boundaries of Ω_h coincide with some element boundaries of the immersed grid γ'_h . If we now choose a space Q'_h which contains piece wise polynomials of degree compatible with that of the elements in V_h , we can avoid the projection operator altogether, and are allowed to write

$$\lambda'_h = -\llbracket \nabla u_h \cdot \mathbf{n} \rrbracket + \beta h^{-1} (u_h - g).$$

Inserting this back into the first equation we get the variational problem (in which no multiplier is involved) of finding $u_h \in V_h$ such that

$$\begin{aligned} (\nabla u_h, \nabla v_h)_\Omega - \langle \llbracket \nabla u_h \cdot \mathbf{n} \rrbracket, v_h \rangle_\gamma - \langle \llbracket \nabla v_h \cdot \mathbf{n} \rrbracket, u_h \rangle_\gamma + \beta \langle h^{-1} u_h, v_h \rangle_\gamma \\ = (f, v_h)_\Omega - \langle \llbracket \nabla v_h \cdot \mathbf{n} \rrbracket, g \rangle_\gamma + \beta \langle h^{-1} g, v_h \rangle_\gamma, \end{aligned} \tag{14}$$

holds for every $v_h \in V_h$.

Equation (14) represents the Nitsche method [27] applied to Problem (1). Owing to the non-matching nature of our discretization, the immersed and background mesh facets are not expected to be aligned in general. If indeed for all facets (elements) K of γ and facets F of Ω_h we have $\mathcal{H}^{d-1}(K \cap F) = 0$, with \mathcal{H}^{d-1} denoting the Hausdorff measure in $d - 1$ -dimensions, then the variational problem can be simplified since all the jump terms vanish.

3.3. The cut-FEM method

When using the cut-FEM discretization approach, one changes the perspective of the original variational problem, which is no longer solved on a single space defined globally on Ω : instead, one solves two separate problems on the two domains ω and $\Omega \setminus \omega$. This approach gives additional flexibility on the type of problems that can be solved. For example, problems with more general transmission conditions across γ where the solution u is allowed to jump. Moreover, separating the problem transforms the cut-FEM method into a boundary-fitted approach, where the approximation space is changed to resolve the interface.

The usual approach to impose constraints on γ in the cut-FEM method is to use Nitsche’s method applied on the two subdomains separately:

$$\begin{cases} \Omega^1 := \omega, \\ \Omega^2 := \Omega \setminus \bar{\omega}. \end{cases}$$

In this context, it is necessary to take special care of those elements of Ω_h that are cut by γ . First, we introduce the corresponding computational meshes Ω_h^i given by

$$\Omega_h^i := \{T \in \Omega_h : T \cap \Omega_i \neq \emptyset\} \quad i = 1, 2,$$

and notice that both meshes share the set of cells intersected by the curve γ , namely

$$\tau := \{T \in \Omega_h : T \cap \gamma \neq \emptyset\}, \quad \tau^i := \{\tilde{T}^i := T \cap \Omega^i, T \in \tau\},$$

where we distinguish between entire cells that intersect γ (these are in the set τ) and cut cells (with arbitrary polytopal shape, which are in the set τ^i), i.e., $(T \in \tau) = (\tilde{T}^1 \in \tau^1) \cup (\tilde{T}^2 \in \tau^2) \cup (\gamma \cap T)$ with $\tilde{T}^1 \cap \tilde{T}^2 = \emptyset$.

When defining a finite element space on these elements, one uses the same definition of the original finite element space defined on entire elements $T \in \tau$, which is then duplicated and restricted to the corresponding domain, i.e., one introduces on Ω_h^i , $i = 1, 2$ the discrete spaces

$$V_h^i := V_h(\Omega_h^i)|_{\Omega^i} = \{v_h|_{\Omega^i}, v_h \in V_h(\Omega_h^i)\}.$$

Then, applying twice Nitsche’s method requires to find $u_h^i \in V_h^i$ such that

$$a_h^i(u_h^i, v_h^i) = l_h^i(v_h^i) \quad \forall v_h^i \in V_h^i \quad i = 1, 2,$$

with

$$\begin{aligned} a_h^i(u_h^i, v_h^i) &= (\nabla u_h^i, \nabla v_h^i)_{\Omega_h^i \cap \Omega^i} - \langle \llbracket \nabla u_h^i \cdot \mathbf{n} \rrbracket, v_h^i \rangle_\gamma - \langle u_h^i, \llbracket \nabla v_h^i \cdot \mathbf{n} \rrbracket \rangle_\gamma \\ &\quad + \frac{\beta_1}{h} \langle u_h^i, v_h^i \rangle_\gamma, \end{aligned} \tag{15}$$

and

$$l_h^i(v_h^i) = (f, v_h^i)_{\Omega_h^i \cap \Omega^i} + \left\langle g, \frac{\beta_1}{h} v_h^i - \llbracket \nabla v_h^i \cdot \mathbf{n} \rrbracket \right\rangle_\gamma. \tag{16}$$

This formulation is known to suffer from the so called *small-cut* problem deriving from the fact that the size of the cuts $T \cap \Omega^i$ cannot be controlled and hence can be arbitrarily small. This may result in a loss of coercivity for the bilinear forms a_h^i when the size of a cut cell goes to zero. As shown in [11], the formulation can be stabilized by adding to the bilinear form the following penalty term acting on the interior or exterior faces of the intersected cells, depending on the domain Ω_h^i :

$$\mathcal{G}_h^i := \{F = \bar{T}_+ \cap \bar{T}_- : T_+ \in \tau, T_- \in \Omega_h^i\} \quad i = 1, 2, \tag{17}$$

$$j_h^i(u, v) := \beta_2 \sum_{F \in \mathcal{G}_h^i} \langle h_F \llbracket \nabla u \cdot \mathbf{n} \rrbracket, \llbracket \nabla v \cdot \mathbf{n} \rrbracket \rangle_F, \tag{18}$$

where β_2 is a positive penalty parameter and h_F the size of $F \in \mathcal{G}_h^i$. With such definition at hand we set

$$V_h := V_h^1 + V_h^2, \text{ with elements } v_h = \begin{cases} v_h^1 & \text{in } \Omega^1, \\ v_h^2 & \text{in } \Omega^2, \end{cases} \tag{19}$$

$$a_h(u_h, v_h) := \sum_{i=1}^2 (a_h^i(u_h^i, v_h^i) + j_h^i(u_h^i, v_h^i)), \tag{20}$$

$$l_h(v_h) := \sum_{i=1}^2 l_h^i(v_h^i), \tag{21}$$

and the method reads: find $u_h \in V_h$ such that

$$a_h(u_h, v_h) = l_h(v_h) \quad \forall v_h \in V_h.$$

The penalization term in (18), usually called *ghost-penalty*, is not the only choice that allows to recover optimal rates. Another practical alternative, developed earlier in [18], consists in associating “bad elements”, i.e. elements with small cuts, with suitable neighboring elements having sufficiently large intersections with the domain in order to extend the polynomial approximation. This technique is usually referred to as *agglomeration* and it has been applied to several unfitted methods in recent years [10,23].

4. Integration of coupling terms

In all three methods, some terms need to be integrated over the non-matching interface γ . For example, in the Lagrange multiplier method,

we need to compute $\langle \lambda_h, v_h \rangle_\gamma$ where $v_h \in V_h$ and $\lambda_h \in Q_h$, while in the Nitsche’s interface penalization method and in the cut-FEM method, one needs to integrate terms of the kind $\langle \beta h^{-1} u_h, v_h \rangle_\gamma$, where both u_h and v_h belong to V_h , but the integral is taken over γ .

We start by focusing our attention on the term $\langle \lambda_h, v_h \rangle_\gamma$ where $v_h \in V_h$ and $\lambda_h \in Q_h$. This is delicate to assemble as it is the product of trial and test functions living on *different* meshes and it encodes the interaction between the two grids. Let $K \in \gamma_h$ and F_K be the map $F_K : \hat{K} \rightarrow K$ from the reference immersed cell \hat{K} to the physical cell K , $JF_K(x)$ the determinant of its Jacobian and assume to have a quadrature rule $\{\hat{x}_q, w_q\}_{q=1}^{N_q}$ on \hat{K} . Since the discrete functions are piecewise polynomials, we can use the scalar product in $L^2(\gamma)$ instead of the duality pairing, then standard finite element assembly reads:

$$\langle \lambda_h, v_h \rangle_\gamma = \sum_{K \in \gamma_h} \langle \lambda_h, v_h \rangle_K. \tag{22}$$

In the forthcoming subsections we discuss three strategies to compute this integral. Identical considerations apply in order to compute the term $\beta \langle h^{-1} u_h, v_h \rangle_\gamma$ in (14).

In general, such integrals are always computed using quadrature formulas. What changes is the algorithm that is used to compute these formulas, and the resulting accuracy. Independently on the strategy that is used to compute the quadrature formulas, all algorithms require the efficient identification of pairs of potentially overlapping cells, or to identify background cells where quadrature points may fall. This task can be performed efficiently by using queries to R-trees of axis-aligned bounding boxes of cells for both the background and immersed mesh.

An R-tree is a data structure commonly used for spatial indexing of multi-dimensional data that relies on organizing objects (e.g., points, lines, polygons, or bounding boxes) in a hierarchical manner based on their spatial extents, such that objects that are close to each other in space are likely to be located near each other in the tree.

In an R-tree, each node corresponds to a rectangular region that encloses a group of objects, and the root node encloses all the objects. Each non-leaf node in the tree has a fixed number of child nodes, and each leaf node contains a fixed number of objects. R-trees support efficient spatial queries such as range queries, nearest neighbor queries, and spatial joins by quickly pruning parts of the tree that do not satisfy the query constraints.

In particular, we build two R-tree data structures to hold the bounding boxes of every cell of both the background and the immersed meshes. Spatial queries are performed traversing the R-tree structure generated by the `Boost.Geometry` library [8]. The construction of an R-tree with M objects has a computational cost that is proportional to $O(M \log(M))$, while the cost of a single query is $O(\log(M))$.

4.1. Integration driven by the immersed mesh

Applying straightforwardly a given quadrature rule defined over γ to equation (22) gives:

$$\langle \lambda_h, v_h \rangle_\gamma \approx \sum_{K \in \gamma_h} \sum_{q=0}^{N_q} \lambda_h(F_K(\hat{x}_q)) v_h(F_K(\hat{x}_q)) JF_K(\hat{x}_q) w_q, \tag{23}$$

with respect to some reference quadrature points \hat{x}_q and weights w_q , letting the immersed domain drive the integration. In this case the computational complexity stems from the evaluation of the terms $v_h(F_K(\hat{x}_q))$, since the position within the background mesh Ω_h of the quadrature point $F_K(\hat{x}_q)$ is not known *a-priori* (see Fig. 2). A possible algorithm for the evaluation of $v_h(F_K(\hat{x}_q))$ can be summarized as follows:

- Compute the physical point $y = F_K(\hat{x}_q)$;
- Find the cell $T \in \Omega_h$ s.t. $y \in T$;
- Given the shape function $\hat{v}_h(\hat{x})$ in the reference element \hat{T} , compute $\hat{v}_h(G_T^{-1}(y))$, where $G_T : \hat{T} \rightarrow T$ denotes the reference map associated to $T \in \Omega_h$ for the background domain.




Fig. 2. Squares: DoF for linear basis functions attached to some elements K_j of γ intersecting a background element $T_i \in \Omega_h$. Dots: DoF for a linear Lagrangian basis attached to cells T_i . Crosses: quadrature points corresponding to a Gaussian quadrature rule of order 3 on elements of γ .

The first part of the computation takes linear time in the number of cells of the immersed mesh γ_h , while the second part requires a computational cost that scales logarithmically with the number of cells of the background grid *for each quadrature point*. Finally, the evaluation of the inverse mapping G_T^{-1} requires Newton-like methods for general unstructured meshes or higher order mappings (i.e., whenever F_K is non-affine). Overall, implementations based on R-tree traversal will result in a computational cost that scales at least as $O((M + N) \log(N))$, where M is the total number of quadrature points (proportional to the number of cells of the immersed mesh), and N is the number of cells of the background mesh.

The main drawbacks of this approach are twofold: i) on one side it does not yield the required accuracy even if quadrature rules with the appropriate order are used, due to the piecewise polynomial nature of the integrands, that may have discontinuities within the element K , and ii) since the couplings between the two grids are based solely on a collection of quadrature points on the immersed surface, it may happen that two elements overlap, but no quadrature points fall within the intersection of the two elements.

This is illustrated in a pathological case in Fig. 3 (top), where we show a zoom-in of a solution computed with an insufficient number of quadrature points, and an overly refined background grid. In this case the quadrature rule behaves like a collection of Dirac delta distributions, and since the resolution of the background grid is much finer than the resolution of the immersed grid, one can recognize in the computed solution the superposition of many small fundamental solutions, that, in the two-dimensional case, behave like many logarithmic functions centered at the quadrature points of γ_h .

These issues with the quadrature driven approach may hinder the convergence properties of the methods, as shown in [5] for a 2D-2D problem, and require a careful equilibrium between the resolution of the immersed grid γ_h , the choice of the quadrature formula on K , and the resolution of the background grid Ω_h . Alternatively, one can follow a different approach that is based on the identification of the intersections between the two grids, and on the use of a quadrature rules defined on the intersection of the two elements, to remove the artifacts discussed above (at the cost of computing the intersection between two non-matching grids), as shown in Fig. 3 (bottom).

4.2. Integration on mesh intersections

An accurate computation of the interface terms may be performed by taking into account the intersection between the two grids. First, the non-empty intersections $\tilde{E} := T \cap K \neq \emptyset$ between any $T \in \Omega_h$ and $K \in \gamma_h$ are identified and the intersection is computed accordingly. Then, given that the restriction of v_h to \tilde{E} is smooth, a suitable quadrature formula

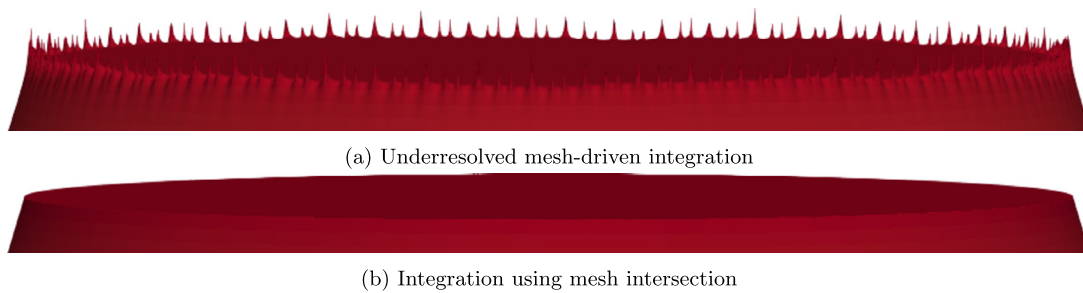


Fig. 3. Comparison between exact and non exact integration (zoomed-in on the interface in order to highlight the kinks) for a pathological case. (a) Spurious kinks around the interface, on the location of the immersed quadrature points. (b) Well-resolved solution around the interface.

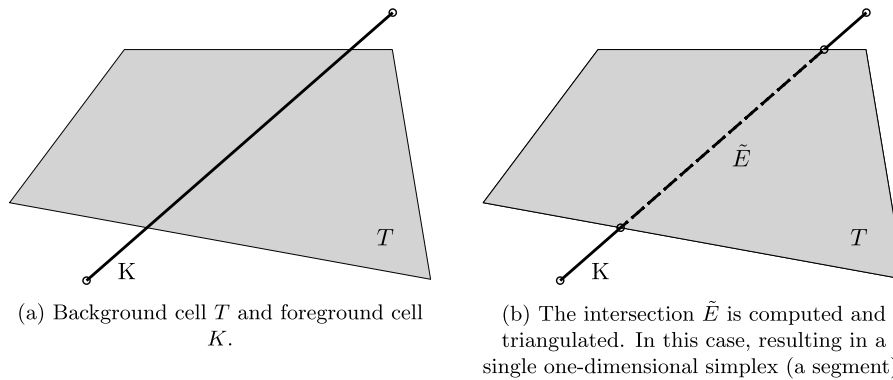


Fig. 4. Triangulation of the intersected region \tilde{E} for a line immersed in 2D. The one-dimensional cell K is allowed to be positioned arbitrarily w.r.t. to the two-dimensional quadrilateral T .

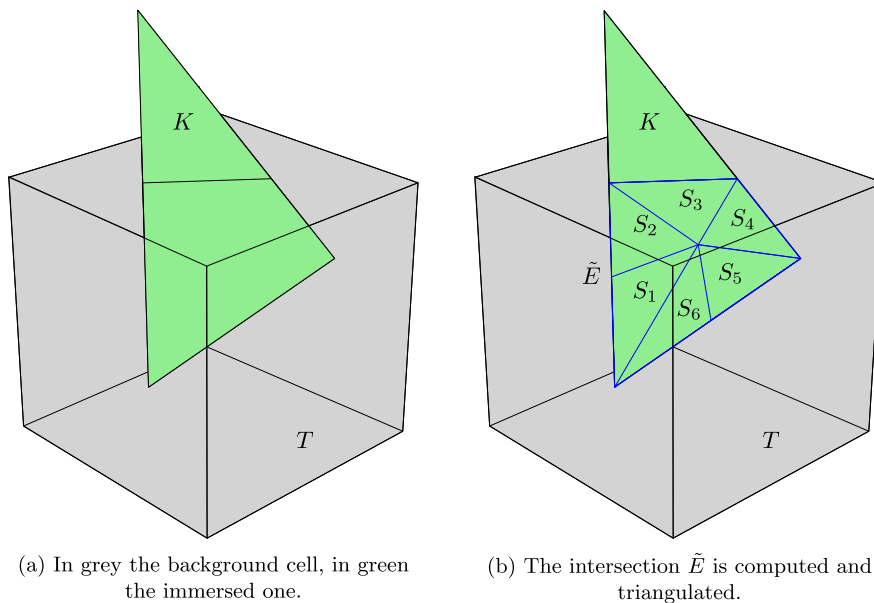


Fig. 5. Triangulation of the intersected region \tilde{E} for a triangle cutting a square and the relative sub-tessellation.

can be applied in \tilde{E} . Since \tilde{E} is a polygon in general, we use a sub-tessellation technique, consisting in splitting \tilde{E} into sub-elements $S \in S_{\tilde{E}}$ such that $\tilde{E} = \bigcup_{S \in S_{\tilde{E}}} S$ and we use standard Gaussian quadrature rules on each of these sub-elements. See Figs. 4 and 5 for an example in two and three dimensions, respectively.

In conclusion, interface terms such as $\langle \lambda_h, v_h \rangle_\gamma$ are assembled by summing the contribution of each intersection \tilde{E} computed by appropriate quadrature:

$$\langle \lambda_h, v_h \rangle_{\tilde{E}} = \sum_{S \in S_{\tilde{E}}} \langle \lambda_h, v_h \rangle_S \approx \sum_{S \in S_{\tilde{E}}} \sum_{q=1}^{N_q} \lambda_h(F_S(\hat{\mathbf{x}}_q)) v_h(F_S(\hat{\mathbf{x}}_q)) J F_S(\hat{\mathbf{x}}_q) w_q, \tag{24}$$

where $F_S : \hat{K} \rightarrow S_S$ is the mapping from the reference element \hat{K} to S and $\{\hat{\mathbf{x}}_q, w_q\}_{q=1}^{N_q}$ a suitable quadrature rule defined on \hat{K} .

The efficient identification of pairs of potentially overlapping cells is performed using queries to R-trees of axis-aligned bounding boxes of




Fig. 6. Quadrature points distributed on O , S and I for a cell T cut by the curve γ .

cells for both the background and immersed mesh. While this allows to record the indices of the entries to be allocated during the assembly procedure and the relative mesh iterators, the actual computation of the geometric intersection \tilde{E} between two elements $T \subset \Omega_h$ and $K \subset \gamma$ is computed using the free function `CGAL::intersection()`. The resulting polytope is sub-tessellated into simplices, i.e., $\tilde{E} = \cup_{i=0}^{N_s} S_i$, even though other techniques for numerical quadrature on polygons may be employed [12].

In Figs. 4 and 5 we show this procedure graphically for cells T and K and provide an example for the corresponding sub-tessellations of the intersection \tilde{E} in two and three dimensions. The resulting sub-tessellations are used to construct the quadrature rules $\{Q_i\}_i$ in (24).

The overall complexity of the algorithm is $O((N + M) \log(N))$, where N and M are the numbers of cells in the background and immersed mesh, respectively. Notice, however, that the complexity of the algorithm should be multiplied by the complexity of the `CGAL::intersection()` function, which is $O(nm)$ for the intersection of two polygons with n and m vertices, respectively. In practice, this is not a problem since the number of vertices of the polygons is typically small, and the complexity of the algorithm is dominated by the complexity of the R-tree queries, but these costs are non-negligible for large grids, and many overlapping cells (see, e.g., the discussion in [5]).

4.3. Integration through level set splitting

If a level set description of the immersed domain is available, this may be used to generate quadrature formulas without explicitly computing the geometric intersection (see [11] for some implementation details). It has also been shown in [29] how to generate quadrature rules on different regions of a cut element using Ψ , identified as

$$O = \{(x, y) \in T : \Psi > 0\},$$

$$S = \{(x, y) \in T : \Psi = 0\},$$

$$I = \{(x, y) \in T : \Psi < 0\},$$

where $\Psi : \mathbb{R}^d \rightarrow \mathbb{R}$ is the level set function determining the immersed domain. A typical configuration including quadrature points for each entity is shown in Fig. 6.

In practical computations, however, the interface γ is not available in terms of a simple analytical description of the level set function. Should one still wish to use quadrature formulas based on a level set, a discrete (possibly approximated) level set function Ψ_h that is zero on γ_h must be provided when γ_h is a triangulated surface. Such level set would also allow a robust partitioning of the background computational mesh into cells that are completely inside ω , cells cut by γ , and cells that are completely inside $\Omega \setminus \omega$.

We propose a simple implementation of a discrete level set function Ψ_h constructed from a triangulated interface γ_h . Point classification (i.e., detecting if a point is inside or outside ω) is performed using a query to the CGAL library, to detect if a point is inside or outside the coarsest simplicial mesh bounded by γ_h (denoted by I_h).

We then define a discrete level set function $\Psi_h(x)$ on top of γ_h as follows:

$$\Psi_h(p) = \begin{cases} -d(p, \gamma) & \text{if } p \in I_h, \\ d(p, \gamma) & \text{if } p \notin I_h, \end{cases} \quad (25)$$

where $d(p, \gamma) := \min_{y \in \gamma} d(p, y)$ is constructed by first finding the closest elements of γ to p using efficient R-tree data structures indexing the cells of the immersed triangulation, and then computing the distance between p and those elements.

This procedure is summarized in Algorithm 1. We validate our approach with a manufactured case by choosing randomly distributed points $\{p_i\}_i$ in the interval $[-1, 1]^2$ and computing the relative error $E_r(p) := \frac{|\Psi(p) - \Psi_h(p)|}{|\Psi(p)|}$ for a level set describing a circle of radius $R = 0.3$, and a corresponding approximated grid γ_h .

In Fig. 7 we report the relative error committed by replacing the exact level set with the discrete one (induced by replacing the exact curve with a triangulated one). The initial discretization of γ is chosen so that the error is below 10^{-6} everywhere, and we can safely neglect it when computing convergence rates of the error computed w.r.t. exact solutions known on the analytical level set.

As long as such geometrical error is not dominating we can observe optimal rates in the numerical experiments for cut-FEM. At the same time, this setting allows for a fair comparison between all the schemes. We expect that other practical implementations would require similar tasks and hence that the observed computational cost is representative.

Algorithm 1: Evaluation of the discrete level set (25) for a given mesh γ .

Input : γ polygonal surface mesh,
 I_h coarse triangulation for the interior of γ ,
 $p \in \Omega$.

Output: $d(p, \gamma)$.

```

1 if  $p \in I_h$  then
2   |  $s \leftarrow -1$ 
3 else
4   |  $s \leftarrow +1$ 
5 end
6 Find  $\{K_i\} \in \gamma_h$  nearest to  $p$ .
7 for  $K \in \{K_i\}_i$  do
8   |  $d_i \leftarrow d(p, K)$ ;
9 end
10 Return  $s \cdot \min_i d_i$ 

```

5. Numerical experiments

Our implementation is based on the C++ finite element library `deal.II` [1,2], providing a dimension independent user interface. The implementation of the Lagrange multiplier and of the Nitsche’s interface penalization methods are adapted from the tutorial programs `step-60` and `step-70` of the `deal.II` library, respectively, while the cut-FEM algorithm is adapted from the tutorial program `step-85`, developed in [31].

As a result of this work, we added support and wrappers for the C++ library CGAL ([32], [14]) into the `deal.II` library [2], in order to perform most of the computational geometry related tasks. Thanks to




Fig. 7. The relative error for the discrete level set describing a disk of radius $R = 0.3$ as a function of the distance from the center $\mathbf{x}_c = (\frac{1}{2}, \frac{1}{2})$.

Table 1
Rates in L^2 and H^1 for a smooth u and $H^{-\frac{1}{2}}$ rates for the Lagrange multiplier method.

Results for $\gamma = \gamma^1$ and smooth solution with Lagrange multiplier							
DoF number	$\ u - u_h\ _{0,\Omega}$	$L^2(\Omega)$ rate	$\ u - u_h\ _{1,\Omega}$	$H^1(\Omega)$ rate	$\ \lambda - \lambda_h\ _{-\frac{1}{2},\gamma}$	$H^{-\frac{1}{2}}(\gamma)$ rate	Iter.
389+32	5.579e-02	-	1.879e+00	-	7.402e-02	-	7
1721+64	1.393e-02	1.87	9.391e-01	0.93	1.042e-02	2.83	7
7217+128	3.479e-03	1.94	4.690e-01	0.97	2.805e-03	1.89	9
29537+256	8.691e-04	1.97	2.343e-01	0.98	6.716e-04	2.06	11
119489+512	2.172e-04	1.98	1.171e-01	0.99	2.078e-04	1.69	11

Table 2
 L^2 and H^1 error rates for $\gamma = \gamma^1$ and a smooth solution with Nitsche.

Results for $\gamma = \gamma^1$ and smooth solution with Nitsche					
DoF number	$\ u - u_h\ _{0,\Omega}$	$L^2(\Omega)$ rate	$\ u - u_h\ _{1,\Omega}$	$H^1(\Omega)$ rate	Iter.
389	5.597e-02	-	1.879e+00	-	1
1721	1.396e-02	1.87	9.391e-01	0.93	11
7217	3.487e-03	1.94	4.690e-01	0.97	11
29537	8.712e-04	1.97	2.343e-01	0.98	12
119489	2.177e-04	1.98	1.171e-01	0.99	12

Table 3
 L^2 and H^1 error rates for $\gamma = \gamma^1$ and a smooth solution with cut-FEM.

Results for $\gamma = \gamma^1$ and smooth solution with cut-FEM					
DoF number	$\ u - u_h\ _{0,\Omega}$	$L^2(\Omega)$ rate	$\ u - u_h\ _{1,\Omega}$	$H^1(\Omega)$ rate	Iter.
329	8.756e-02	-	2.1335e+00	-	2
1161	2.2064e-02	1.99	1.0582e+00	1.01	2
4377	5.4875e-03	2.01	5.2351e-01	1.02	15
16953	1.3554e-03	2.02	2.5863e-01	1.02	17
66665	3.3152e-04	2.03	1.2855e-01	1.01	18

the so called *exact computation paradigm* provided by CGAL, which relies on computing with numbers of arbitrary precision, our intersection routines are guaranteed to be robust.

We assume that the background mesh Ω_h is a d -dimensional triangulation and the immersed mesh γ_h is $(d - 1)$ -dimensional with $d = 2, 3$. We validate our implementations with several experiments varying mesh configurations, algorithms, and boundary conditions. The source code used to reproduce the numerical experiments is available from GitHub.³

³ https://github.com/fdrmc/non_matching_test_suite.git.

The tests are designed to analyze the performance of the methods presented in Section 3 in different settings, varying the complexity of the interface and the smoothness of the exact solution in both two and three dimensions. All tests are performed using background meshes made of quadrilaterals or hexahedra and immersed boundary meshes made of segments and quadrilaterals. For the Lagrange multiplier and Nitsche’s interface penalization methods, we perform an initial pre-processing of the background grid Ω by applying a localized refinement around the interface (where most of the error is concentrated), so that the resulting number of degrees of freedom for the variable u_h is roughly the same for all methods. Sample grids resulting from this process are shown in Fig. 8, where the interface has been resolved from Ω_h . We then proceed by computing errors and convergence rates against a manufactured solution under simultaneous refinement of both the background and immersed mesh.

Classical Q^1 Lagrangian elements are used for the background space while piecewise constant elements are used to discretize the Lagrange multiplier. For the Nitsche penalization method (14), we set the penalty parameter as $\beta = 10$. Errors in the H^1 - and L^2 -norm are reported for the main variable while for the Lagrange multiplier we use the discrete norm in (7) which, as already observed, is equivalent to the $H^{-1/2}(\gamma)$ norm on a quasi-uniform mesh. In the case of the Lagrange multiplier method, we report the sum of the number of Degrees of Freedom (DoF) for u_h and λ_h to underline the fact that a larger system must be solved, while rates are computed against the number of DoF of each unknown.

For the (symmetric) penalized methods, the resulting linear systems are solved using a preconditioned conjugate gradient method, with an algebraic multigrid preconditioner based on the Trilinos ML implementation [21], while for the Lagrange multiplier we exploit the preconditioner described in Section 3.1 with the same preconditioned conjugate gradient method for inner solves of the stiffness matrices, and flexible GMRES as the outer solver. In the case of the Lagrange multiplier method, we list the total number of inner iterations required to invert the Schur complement.

Table 4
 L^2 and H^1 error rates for $\gamma = \gamma^2$ and a smooth solution with Lagrange multiplier method.

Results for $\gamma = \gamma^2$ and smooth solution with Lagrange multiplier							
DoF number	$\ u - u_h\ _{0,\Omega}$	$L^2(\Omega)$ rate	$\ u - u_h\ _{1,\Omega}$	$H^1(\Omega)$ rate	$\ \lambda - \lambda_h\ _{-\frac{1}{2},\gamma}$	$H^{-\frac{1}{2}}(\gamma)$ rate	Iter.
1958+256	5.327e-02	-	1.781e+00	-	6.551e-03	-	17
8783+512	1.331e-02	1.85	8.916e-01	0.92	1.183e-03	2.47	20
37037+1024	3.326e-03	1.93	4.458e-01	0.96	5.237e-04	1.18	33
151961+2048	8.312e-04	1.96	2.228e-01	0.98	1.959e-04	1.42	49
615473+4096	2.078e-04	1.98	1.114e-01	0.99	5.327e-05	1.88	53

Table 5
 L^2 and H^1 error rates for $\gamma = \gamma^2$ and a smooth solution with Nitsche.

Results for $\gamma = \gamma^2$ and smooth solution with Nitsche					
DoF number	$\ u - u_h\ _{0,\Omega}$	$L^2(\Omega)$ rate	$\ u - u_h\ _{1,\Omega}$	$H^1(\Omega)$ rate	Iter.
1958	5.327e-02	-	1.781e+00	-	13
8783	1.331e-02	1.85	8.916e-01	0.92	12
37037	3.326e-03	1.93	4.458e-01	0.96	13
151961	8.312e-04	1.96	2.228e-01	0.98	12
615473	2.078e-04	1.98	1.114e-01	0.99	14

Table 6
 L^2 and H^1 error rates for $\gamma = \gamma^2$ and a smooth solution with cut-FEM.

Results for $\gamma = \gamma^2$ and smooth solution with cut-FEM					
DoF number	$\ u - u_h\ _{0,\Omega}$	$L^2(\Omega)$ rate	$\ u - u_h\ _{1,\Omega}$	$H^1(\Omega)$ rate	Iter.
1209	2.2306e-02	-	1.0422e+00	-	2
4487	5.7207e-03	1.96	5.2424e-01	0.99	17
17163	1.4123e-03	2.02	2.5901e-01	1.02	17
67089	3.4377e-04	2.04	1.2865e-01	1.01	19
265249	8.5223e-05	2.01	6.4118e-02	1.00	20

5.1. 2D numerical tests

In order to make the comparison between the three methods as close to real use cases as possible, we do not exploit any *a-priori* knowledge of analytical level set descriptions of the exact interfaces. Indeed, using this information one could expect faster computations of the intersections, and an overall reduction of the computational costs of the assembly routines. Instead, we fix the same discretization of the interface as input data of the computational problem for all three methods.

In particular, we consider as computational domain $\Omega = [-1, 1]^2$ with immersed domains of different shapes originating from an unfitted discretization of two different curves:

- *circle interface*: we let

$$\gamma^1 := \partial B_R(c),$$

- *flower-shaped interface*: we let

$$\gamma^2 := \left\{ (x, y) : \sqrt{x^2 + y^2} - r \left(1 - 2 \frac{y^2}{x^2 + y^2} \right) \left(1 - 16 \frac{x^2 y^2}{(x^2 + y^2)^2} \right) - R = 0 \right\},$$

where the first is a circle of radius R centered at c and the second is a flower-like interface. We choose as parameter values $c = (\frac{1}{2}, \frac{1}{2})$, $R = 0.3$ and $r = 0.1$. The initial discretization of the immersed domains is chosen so that the geometrical error is negligible w.r.t. the discretization errors (see Fig. 8). In the case of the circle interface γ^1 , the discrete interface is generated using a built-in mesh generators from `deal.II` while the flower-shaped interface γ^2 is imported from an external file. When required, we implement a discrete level set function Ψ_h as described in Algorithm 1. Fig. 8 shows a representation of the two interfaces.

In the first two tests, we set up the problem using the method of manufactured solutions, imposing the data f in Ω and the boundary conditions on Γ and γ according to the exact solution

$$u(x, y) := \sin(2\pi x) \sin(2\pi y). \tag{26}$$

Hence, the right hand side of (1) is $f = 8\pi^2 u(x, y)$ and the data on the outer boundary Γ and on γ are computed accordingly. Notice that the smoothness of the solution implies $\lambda \equiv 0$ in this case, meaning that the Dirichlet condition does not constraint the solution u at the interface.

This test is meant to assess the basic correctness of the implementation of the three methods, and corresponds to a case in which the interface is truly not an interface. Instead, when we impose an arbitrary value for the solution at the interface, we expect the solution u to be only in $H^{\frac{3}{2}-\epsilon}(\Omega)$ for any $\epsilon > 0$, even though its local regularity on the two subdomains ω and $\Omega \setminus \omega$ may be higher. This is due to the fact that the gradient of the solution is not a continuous function across the interface, and therefore the solution cannot be in $H^2(\Omega)$. In this case, we cannot expect non-matching methods that does not resolve the interface *exactly*, such as the Lagrange multiplier and the Nitsche’s interface penalization method, to be able to recover the optimal rate of convergence.

In the tables below we also report the number of iterations required in the solution phase in the column ‘Iter.’. We observe for all experiments a similar number of iterations for the three methods (which are independent on the number of degrees of freedom, indicating a good choice of preconditioner for all three methods) even though the solution of the linear system stemming from the Lagrange multiplier method is generally more expensive compared to the other two methods, owing to the higher computational complexity of the preconditioner for the saddle-point problem. The balance in the computational cost of the three different methods is discussed in details in Section 5.3.

5.1.1. Test 1: smooth solution over circular interface

We report in Tables 1, 2, and 3 the errors and computed rates for the Lagrange multiplier, Nitsche’s interface penalization, and cut-FEM method, respectively. In each case, the background variable converges linearly and quadratically in the H^1 - and L^2 -norm, respectively. As for the Lagrange multiplier, we observe a convergence rate close to two instead of the theoretical rate of one, most likely due to the very special exact solution that the multiplier converges to (i.e., the zero function). For a direct comparison, we also report in Fig. 9 (left) the convergence history of all three methods against the number of DoF. These results clearly indicate that for smooth problems with relatively simple interfaces the three methods perform similarly.

5.1.2. Test 2: smooth solution over flower-shaped interface

We report in Tables (4), (5) and (6) the error with the computed rates of convergence for the three schemes. Again we observe the theoretical rates of convergence for all three methods. This time, however, the direct comparison of the errors shown in Fig. 9 (right) indicates that the cut-FEM approach has an advantage over the other two methods. This may be due to the worst approximation properties of the non-matching methods based on Lagrange multipliers and Nitsche’s interface penalization in the presence of high curvature sections of the immersed boundary, as can be seen by comparing the meshes shown in Fig. 8. Despite the fact that the reached accuracy is the same for the Lagrange multiplier, we further note that the Nitsche’s interface penalization method requires less iterations than the Lagrange multiplier method.




Fig. 8. Zoom on pre-processed background grid Ω_h for the circle interface $\gamma = \gamma^1$ (left) and the flower-shaped interface $\gamma = \gamma^2$ (right).




Fig. 9. L^2 , H^1 , and $H^{-\frac{1}{2}}$ errors versus the number of DoF for all schemes applied to γ^1 (left) and to γ^2 (right).

5.1.3. Test 3: non-smooth solution over circular interface

In this test, we fix once more $\gamma = \gamma^1$ and we define an exact solution with a non-zero jump of the normal gradient $\nabla u \cdot \mathbf{n}$ across γ taken from [20], namely

$$u(x, y) = \begin{cases} -\ln(R) & \text{if } |r| \leq R, \\ -\ln(r) & \text{if } |r| > R, \end{cases} \quad (27)$$

where $r := x - c$, implying as right hand side $f = 0$. The Lagrange multiplier associated to this solution is $\lambda(\mathbf{x}) \equiv \lambda = -\frac{1}{R}$, as u solves the following classical interface problem:

$$\begin{cases} -\Delta u & = 0 & \text{in } \Omega \setminus \gamma, \\ u & = -\ln(r) & \text{on } \Gamma, \\ \llbracket \nabla u \cdot \mathbf{n} \rrbracket & = \frac{1}{R} & \text{on } \gamma, \\ \llbracket u \rrbracket & = 0 & \text{on } \gamma. \end{cases} \quad (28)$$

Since the global regularity of the solution is $H^{\frac{3}{2}-\epsilon}(\Omega)$ for any $\epsilon > 0$, theoretically we would expect the convergence rates of the $L^2(\Omega)$, $H^1(\Omega)$, and $H^{-\frac{1}{2}}(\gamma)$ norms of the errors to be 1, 0.5, and 0.5 for the Lagrange multiplier method, the same for the variable u_h in the Nitsche’s interface penalization method, and the optimal convergence rates observed in the smooth case in the case of the cut-FEM method. These are shown in Tables 7, 8 and 9 and plotted in Fig. 10.

The results show a clear advantage in terms of convergence rates and absolute values of the errors for the cut-FEM method. In all cases (both smooth and non-smooth), the Nitsche’s interface penalization method and the Lagrange multiplier give essentially the same errors. A simple way to improve the situation for the latter two methods would be to use a weighted norm during the computation of the error, which was proven to be localized at the interface in [20]. This would allow to reduce the overall error, but it would still result in a solution that does not capture correctly the jump of the normal gradient across the interface, which is the main source of error.

5.2. 3D numerical tests

We mimic the tests reported above for the two-dimensional setting also in the three-dimensional case, but we restrict our analysis to the case of a sphere immersed in a box. We fix $\Omega = [-1, 1]^3$ and consider as immersed interface a sufficiently fine discretization of the sphere $\gamma^3 := \partial B_R(c)$ with $c = (\frac{1}{2}, \frac{1}{2}, \frac{1}{2})$, $R = 0.3$. In this setting, we consider two test cases with smooth and non-smooth solution.

5.2.1. Test 1: smooth solution over spherical interface

We proceed analogously to the previous section by computing convergence rates when the solution u is smooth and defined as

$$u(x, y, z) := \sin(2\pi x) \sin(2\pi y) \sin(2\pi z), \quad (29)$$

Table 7
 L^2 -error and H^1 -error for non smooth u in (27) and for the multiplier.

Results for $\gamma = \gamma^1$ and non-smooth solution with Lagrange multipliers							
DoF number	$\ u - u_h\ _{0,\Omega}$	$L^2(\Omega)$ rate	$\ u - u_h\ _{1,\Omega}$	$H^1(\Omega)$ rate	$\ \lambda - \lambda_h\ _{-\frac{1}{2},\gamma}$	$H^{-\frac{1}{2}}(\gamma)$ rate	Iter.
389+32	1.413e-02	-	3.656e-01	-	3.354e-02	-	2
1721+64	5.698e-03	1.22	2.254e-01	0.65	1.112e-02	1.59	2
7217+128	3.248e-03	0.78	1.621e-01	0.46	6.273e-03	0.83	2
29537+256	1.796e-03	0.84	1.141e-01	0.50	4.094e-03	0.62	2
119489+512	1.099e-03	0.70	8.015e-02	0.51	2.845e-03	0.52	2




Fig. 10. L^2 , H^1 , and $H^{-\frac{1}{2}}$ errors versus the number of DoF for all schemes applied to γ^1 , with non-smooth solution.

Table 8
 L^2 -error, H^1 -error for non smooth u in (27).

Results for non-smooth u with Nitsche					
DoF number	$\ u - u_h\ _{0,\Omega}$	$L^2(\Omega)$ rate	$\ u - u_h\ _{1,\Omega}$	$H^1(\Omega)$ rate	Iter.
389	9.216e-03	-	3.667e-01	-	1
1721	3.324e-03	1.37	2.286e-01	0.64	11
7217	1.936e-03	0.75	1.640e-01	0.46	11
29537	9.957e-04	0.94	1.151e-01	0.50	12
119489	5.016e-04	0.98	8.070e-02	0.51	13

Table 9
 L^2 -error, H^1 -error for non smooth u in (27) with cut-FEM.

Results for non-smooth u with cut-FEM					
DoF number	$\ u - u_h\ _{0,\Omega}$	$L^2(\Omega)$ rate	$\ u - u_h\ _{1,\Omega}$	$H^1(\Omega)$ rate	Iter.
329	3.1670e-02	-	5.6660e-01	-	2
1161	1.5998e-03	4.31	1.2352e-01	2.20	2
4377	3.6791e-04	2.12	5.8327e-02	1.08	16
16953	8.5282e-05	2.11	2.7307e-02	1.09	18
66665	2.1914e-05	1.96	1.3726e-02	0.99	19

which corresponds to the right hand side $f = 12\pi^2 u(x, y, z)$. We report in Tables 10, 11 and 12 the error rates for the three schemes, while in Fig. 11 (left) we plot errors against the number of DoF.

The convergence rates are as expected and the results are in line with the two-dimensional case. The Lagrange multiplier method and the interface penalization method give again very close computational errors for u_h . As in the smooth two-dimensional case, this test should

only be considered as a validation of the code and of the error computation, since the interface does not have any effect on the computational solutions.

5.2.2. Test 2: non-smooth solution over spherical interface

We consider the test case in [20]:

$$u(x, y) = \begin{cases} \frac{1}{R} & \text{if } |r| \leq R, \\ \frac{1}{|r|} & \text{if } |r| > R, \end{cases} \tag{30}$$

where $r := x - c$, $f = 0$, and $R = 0.3$. Analogously to the two dimensional case, the multiplier associated to this problem is $\lambda(x) = \lambda = -\frac{1}{R^2}$, since u solves the following problem:

$$\begin{cases} -\Delta u & = 0 & \text{in } \Omega \setminus \gamma, \\ u & = \frac{1}{|r|} & \text{on } \Gamma, \\ \llbracket \nabla u \cdot \mathbf{n} \rrbracket & = \frac{1}{R^2} & \text{on } \gamma, \\ \llbracket u \rrbracket & = 0 & \text{on } \gamma. \end{cases} \tag{31}$$

We report in Tables 13, 14, 15 the error rates for the three schemes and in Fig. 11 (right) the error decay in $L^2(\Omega)$ and $H^1(\Omega)$ for u_h and the decay in $H^{-\frac{1}{2}}(\gamma)$ for λ_h . A contour plot of the discrete solution u_h obtained with Nitsche’s penalization method is shown in Fig. 12.

The difference between the three methods is less evident in terms of absolute values of the errors when compared to the two dimensional case: while it is still clear that the convergence rates of cut-FEM are higher compared to the other two methods, the difference between the three methods is smaller. In particular, the Nitsche’s interface penalization method seems to perform better than the Lagrange multiplier method when considering the L^2 norm.

Table 10
 L^2 and H^1 error rates for $\gamma = \gamma^3$ and a smooth solution with Lagrange multiplier.

Results for $\gamma = \gamma^3$ and smooth solution with Lagrange multiplier							
DoF number	$\ u - u_h\ _{0,\Omega}$	$L^2(\Omega)$ rate	$\ u - u_h\ _{1,\Omega}$	$H^1(\Omega)$ rate	$\ \lambda - \lambda_h\ _{-\frac{1}{2},\gamma}$	$H^{-\frac{1}{2}}(\gamma)$ rate	Iter.
4127+96	6.416e-02	-	2.432e+00	-	0.1027	-	9
37031+384	1.590e-02	1.91	1.212e+00	0.95	0.0287	1.84	14
313007+1536	3.963e-03	1.95	6.051e-01	0.98	0.0069	2.06	19
2572511+6144	9.996e-04	1.96	3.024e-01	0.99	0.0019	1.89	14




Fig. 11. L^2 , H^1 , and $H^{-\frac{1}{2}}$ error versus the number of DoF for all schemes applied to γ^3 with a smooth solution u (left) and with a non-smooth solution u (right).




Fig. 12. Background mesh Ω_h and immersed mesh of the sphere interface γ_h for the three dimensional case (left) and section of the contour plot for the approximate solution u_h in (30), $\gamma = \gamma^3$.

Table 11
 L^2 and H^1 error rates for $\gamma = \gamma^3$ and a smooth solution with Nitsche.

Results for $\gamma = \gamma^3$ and smooth solution with Nitsche					
DoF number	$\ u - u_h\ _{0,\Omega}$	$L^2(\Omega)$ rate	$\ u - u_h\ _{1,\Omega}$	$H^1(\Omega)$ rate	Iter.
4127	6.411e-02	-	2.432e+00	-	17
37031	1.588e-02	2.01	1.212e+00	1.00	14
313007	3.959e-03	2.00	6.050e-01	1.00	14
2572511	9.887e-04	2.00	3.023e-01	1.00	14

Table 12
 L^2 and H^1 error rates for $\gamma = \gamma^3$ and a smooth solution with cut-FEM.

Results for $\gamma = \gamma^3$ and smooth solution with cut-FEM					
DoF number	$\ u - u_h\ _{0,\Omega}$	$L^2(\Omega)$ rate	$\ u - u_h\ _{1,\Omega}$	$H^1(\Omega)$ rate	Iter.
5163	7.0847e-02	-	2.4835e+00	-	24
36781	1.7743e-02	2.12	1.2479e+00	1.05	21
278157	4.5124e-03	2.03	6.2396e-01	1.03	22
2160541	1.1302e-03	2.03	3.1154e-01	1.02	22

5.3. Computational times

In order to better understand the performance of the three methods, we consider a breakdown of the computational costs into work precision diagrams. These provide a more fair measure of efficiency as they take into account the computational cost required to reach a given accuracy.

We report hereafter a breakdown of the computational times needed by our implementations of the three proposed methods. All computations were carried out on a 2.60 GHz Intel Xeon processor. For each benchmark we report the average time required to solve 10 times the 3D smooth Problem 5.2.1. We recall that a smooth solution implies a vanishing Lagrange multiplier λ , hence the solution u is not constrained. We compute separately the required CPU times (in seconds) for the main tasks that each scheme has to perform. On a quasi-uniform mesh, the number N of background cells in Ω_h scales with $\mathcal{O}(h_\Omega^{-3})$, and we expect the assembly of the stiffness matrix to scale linearly in the number of cells. On the other hand, the number of facets in γ_h scales with $\mathcal{O}(h_\gamma^{-2})$; in our experiments, the ratio h_Ω/h_γ is kept fixed, therefore we expect the assembly of the coupling terms $\langle \lambda, v \rangle_\gamma$ and $\langle u, v \rangle_\gamma$ to scale with $\mathcal{O}(N^{\frac{2}{3}})$.

This is indeed what we observe in the experiments as shown in CPU breakdown plots of Fig. 13 for each method.

The three schemes have comparable computational times. In particular, the Nitsche interface penalization method exhibits lower global assembly times compared to the others. However, it is well known that the number of iterations required to solve the algebraic problem is influenced by the choice of the penalty parameter β in (14), which determines simultaneously also the accuracy of the numerical solution u_h . This can be better seen in work-precision diagrams, where we compare the CPU times to solve each refinement cycle versus the L^2 error for both test problems (29) and (30) in Fig. 14. In the smooth scenario, results for the Lagrange multiplier and Nitsche's interface penalization methods are almost overlapping both in terms of time and accuracy, while cut-FEM shows larger computational times. The situation is different in the non-smooth case where cut-FEM better captures the discontinuity at the interface and thus gives more accurate results, with a larger cost in terms of time for low degrees of freedom count, and with smaller cost for large degrees of freedom count, owing to the better convergence rate of the method. These results indicate that the additional implementation complexity eventually pays back for non-smooth solutions. This would be even more the case with higher order elements, as cut-FEM would keep optimal rates while the other methods would not be optimal.

Based on the higher efficiency of the Lagrange multiplier and Nitsche penalization method in the smooth case, we speculate that these methods can be made competitive also in the non-smooth case by an improved local refinement strategy [19].

6. Conclusions

The numerical solution of partial differential equations modeling the interaction of physical phenomena across interfaces with complex, possibly moving, shapes is of great importance in many scientific fields. Boundary unfitted methods offer a valid alternative to remeshing and Arbitrary Lagrangian Eulerian formulations, but require care in the representation of the coupling terms between the interface and the bulk equations.

In this work, we performed a comparative analysis of three non-matching methods, namely the Lagrange multiplier method (or fictitious domain method), the Nitsche's interface penalization method, and the cut-FEM method, in terms of accuracy, computational cost, and implementation effort.

We presented the major algorithms used to integrate coupling terms on non-matching interfaces, discussed the benefit of computing accurate quadrature rules on mesh intersections, and concluded our analysis with a set of numerical experiments in two and three dimensions.

Our results show that accurate quadrature rules can significantly improve the accuracy of the numerical methods, and that there are cases in which simpler methods, like the Nitsche's interface penalization method, are competitive in terms of accuracy per computational effort. In general, the additional implementation burden of the cut-FEM method is justified by the higher accuracy that it achieves, especially in three dimensions and for the solution of non-smooth problems.

We believe that this paper provides a valuable resource for researchers and practitioners working on numerical methods for interface problems, particularly in the context of elliptic PDEs coupled across heterogeneous dimensions. Source codes used to produce the results of this paper are available at github.com/fdrmc/non_matching_test_suite.git, and on the platform Code Ocean codeocean.com/capsule/1296846.

Data availability

No data was used for the research described in the article.

Acknowledgements

The author Lucia Gastaldi is member of INdAM Research group GNCS and she is partially supported by PRIN/MUR on grant No. 20227K44ME. The author Luca Heltai was partially supported by PRIN/MUR on grant No. 2022WKWZA8.

References

- [1] D. Arndt, W. Bangerth, D. Davydov, T. Heister, L. Heltai, M. Kronbichler, M. Maier, J.-P. Pelteret, B. Turcksin, D. Wells, The deal.II finite element library: design, features, and insights, *Comput. Math. Appl.* 81 (Jan. 2021) 407–422.
- [2] D. Arndt, W. Bangerth, M. Feder, M. Fehling, R. Gassmüller, T. Heister, L. Heltai, M. Kronbichler, M. Maier, P. Munch, J.-P. Pelteret, S. Stöck, B. Turcksin, D. Wells, The deal.II library, version 9.4, *J. Numer. Math.* (2022).
- [3] I. Babuška, The finite element method with Lagrangian multipliers, *Numer. Math.* 20 (3) (1973) 179–192.
- [4] D. Boffi, F. Brezzi, M. Fortin, *Mixed Finite Element Methods and Applications*, Springer Series in Computational Mathematics, vol. 44, Springer, Heidelberg, 2013.
- [5] D. Boffi, F. Credali, L. Gastaldi, On the interface matrix for fluid-structure interaction problems with fictitious domain approach, *Comput. Methods Appl. Mech. Eng.* 401 (part B) (2022) 115650.
- [6] D. Boffi, L. Gastaldi, A fictitious domain approach with Lagrange multiplier for fluid-structure interactions, *Numer. Math.* 135 (3) (2017) 711–732.
- [7] D. Boffi, L. Gastaldi, L. Heltai, C.S. Peskin, On the hyper-elastic formulation of the immersed boundary method, *Comput. Methods Appl. Mech. Eng.* 197 (25–28) (2008) 2210–2231.
- [8] Boost, Boost C++ Libraries, <http://www.boost.org/>, 2015.

Table 13
 L^2 and H^1 error rates for $\gamma = \gamma^3$ and non-smooth solution u in (30) with Lagrange multiplier.

Results for $\gamma = \gamma^3$ and non-smooth solution with Lagrange multiplier							
DoF number	$\ u - u_h\ _{0,\Omega}$	$L^2(\Omega)$ rate	$\ u - u_h\ _{1,\Omega}$	$H^1(\Omega)$ rate	$\ \lambda - \lambda_h\ _{-\frac{1}{2},\Gamma}$	$H^{-\frac{1}{2}}(\gamma)$ rate	Iter.
4127+96	3.907e-02	-	1.251e+00	-	0.4552	-	2
37031+384	2.041e-02	0.89	8.119e-01	0.59	0.2598	0.81	3
313007+1536	1.095e-02	0.88	5.764e-01	0.48	0.1806	0.52	3
2572511+6144	6.214e-03	0.81	4.253e-01	0.43	0.1216	0.57	3

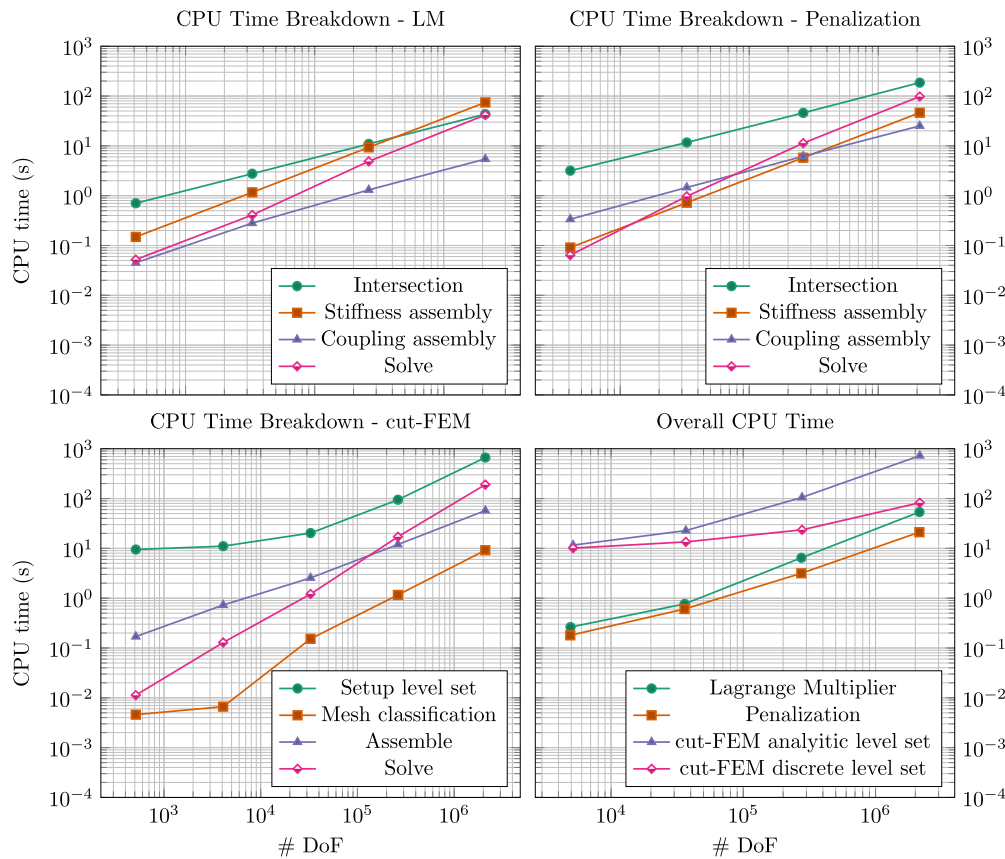


Fig. 13. Breakdown of CPU times for the three schemes. In (bottom left) *Setup level set* indicates the time required to interpolate the discrete level set described in Algorithm 1 onto the finite element space. *Mesh classification* shows the time needed to partition the computational mesh and classify the cells in *cut*, *interior* or *outside* cells. Bottom right: overall CPU times to assemble the algebraic system for the smooth 3D test. For the cut-FEM method, we also show the CPU time obtained by using an analytical representation of the interface through an analytic level set function.

Table 14
 Rates for $\gamma = \gamma^3$ and a non-smooth solution u in (30) with Nitsche.

Results for $\gamma = \gamma^3$ and non-smooth solution with Nitsche					
DoF number	$\ u - u_h\ _{0,\Omega}$	$L^2(\Omega)$ rate	$\ u - u_h\ _{1,\Omega}$	$H^1(\Omega)$ rate	Iter.
4127	8.688e-02	-	1.562e+00	-	19
37031	2.157e-02	2.01	8.885e-01	0.81	16
313007	4.893e-03	2.14	5.724e-01	0.63	16
2572511	1.887e-03	1.37	3.954e-01	0.53	16

Table 15
 L^2 and H^1 error for $\gamma = \gamma^3$ and a non-smooth solution u in (30) with cut-FEM.

Results for $\gamma = \gamma^3$ and non-smooth solution with cut-FEM					
DoF number	$\ u - u_h\ _{0,\Omega}$	$L^2(\Omega)$ rate	$\ u - u_h\ _{1,\Omega}$	$H^1(\Omega)$ rate	Iter.
5163	6.3609e-02	-	1.3046e+00	-	25
36781	6.9076e-03	3.39	4.5097e-01	1.62	22
278157	1.5076e-03	2.26	2.2098e-01	1.06	24
2160541	3.3802e-04	2.19	9.3590e-02	1.26	24

[9] S.C. Brenner, L.R. Scott, *The Mathematical Theory of Finite Element Methods*, Texts in Applied Mathematics, vol. 15, Springer, 2008.

[10] E. Burman, M. Cicuttin, G. Delay, A. Ern, An unfitted hybrid high-order method with cell agglomeration for elliptic interface problems, *SIAM J. Sci. Comput.* 43 (2) (2021) A859–A882.

[11] E. Burman, S. Claus, P. Hansbo, M.G. Larson, A. Massing, Cutfem, *Discretizing geometry and partial differential equations*, *Int. J. Numer. Methods Eng.* 104 (7) (2015) 472–501.

[12] A. Cangiani, E. Georgoulis, P. Houston, *hp*-version discontinuous Galerkin methods on polygonal and polyhedral meshes, *Math. Models Methods Appl. Sci.* 24 (2014) 05.

[13] J. Donea, S. Giuliani, J. Halleux, An arbitrary Lagrangian-Eulerian finite element method for transient dynamic fluid-structure interactions, *Comput. Methods Appl. Mech. Eng.* 33 (1–3) (Sept. 1982) 689–723.

[14] A. Fabri, G.-J. Giezeman, L. Kettner, S. Schirra, S. Schönherr, The CGAL kernel: a basis for geometric computation, *Lect. Notes Comput. Sci.* (1996) 191–202.

[15] V. Girault, R. Glowinski, Error analysis of a fictitious domain method applied to a Dirichlet problem, *Jpn. J. Ind. Appl. Math.* 12 (3) (1995) 487–514.

[16] R. Glowinski, T.-W. Pan, T. Hesla, D. Joseph, A distributed Lagrange multiplier/fictitious domain method for particulate flows, *Int. J. Multiph. Flow* 25 (5) (1999) 755–794.

[17] R. Glowinski, T.-W. Pan, J. Périaux, A fictitious domain method for Dirichlet problem and applications, *Comput. Methods Appl. Mech. Eng.* 111 (3–4) (1994) 283–303.




Fig. 14. Work-precision diagrams for the two 3D tests.

[18] J. Haslinger, Y. Renard, A new fictitious domain approach inspired by the extended finite element method, *SIAM J. Numer. Anal.* 47 (2) (2009) 1474–1499.

[19] L. Heltai, W. Lei, Adaptive finite element approximations for elliptic problems using regularized forcing data, *SIAM J. Numer. Anal.* 61 (2) (Mar. 2023) 431–456.

[20] L. Heltai, N. Rotundo, Error estimates in weighted Sobolev norms for finite element immersed interface methods, *Comput. Math. Appl.* 78 (11) (2019) 3586–3604.

[21] M.A. Heroux, R.A. Bartlett, V.E. Howle, R.J. Hoekstra, J.J. Hu, T.G. Kolda, R.B. Lehoucq, K.R. Long, R.P. Pawlowski, E.T. Phipps, A.G. Salinger, H.K. Thornquist, R.S. Tuminaro, J.M. Willenbring, A. Williams, K.S. Stanley, An overview of the trillions project, *ACM Trans. Math. Softw.* 31 (3) (sep 2005) 397–423.

[22] C.W. Hirt, A.A. Amsden, J.L. Cook, An arbitrary Lagrangian-Eulerian computing method for all flow speeds, *J. Comput. Phys.* 135 (1997) 203–216.

[23] A. Johansson, M. Larson, A high order discontinuous Galerkin Nitsche method for elliptic problems with fictitious boundary, *Numer. Math.* 123 (2013) 04.

[24] R. Krause, P. Zulian, A parallel approach to the variational transfer of discrete fields between arbitrarily distributed unstructured finite element meshes, *SIAM J. Sci. Comput.* 38 (C307–C333) (2016) 01.

[25] A. Massing, M.G. Larson, A. Logg, Efficient implementation of finite element methods on nonmatching and overlapping meshes in three dimensions, *SIAM J. Sci. Comput.* 35 (1) (2013) C23–C47.

[26] N. Moës, J. Dolbow, T. Belytschko, A finite element method for crack growth without remeshing, *Int. J. Numer. Methods Eng.* 46 (1) (1999) 131–150.

[27] J. Nitsche, Über ein Variationsprinzip zur Lösung von Dirichlet-Problemen bei Verwendung von Teilräumen, die Keinen Randbedingungen Unterworfen Sind, *Abh. Math. Semin. Univ. Hamb.* 36 (1971) 9–15.

[28] C.S. Peskin, The immersed boundary method, *Acta Numer.* 11 (2002) 479–517.

[29] R.I. Saye, High-order quadrature methods for implicitly defined surfaces and volumes in hyperrectangles, *SIAM J. Sci. Comput.* 37 (2) (2015) A993–A1019.

[30] R. Stenberg, On some techniques for approximating boundary conditions in the finite element method, *J. Comput. Appl. Math.* 63 (1) (1995) 139–148, *Proceedings of the International Symposium on Mathematical Modelling and Computational Methods Modelling* 94.

[31] S. Stiecko, G. Kreiss, A stabilized Nitsche cut element method for the wave equation, *Comput. Methods Appl. Mech. Eng.* 309 (Sept. 2016) 364–387.

[32] The CGAL Project, *CGAL User and Reference Manual*, 5.5 edition, CGAL Editorial Board, 2022.