

Stability and conditioning of immersed finite element methods analysis and remedies

de Prenter, Frits; Verhoosel, Clemens ; Brummelen, Harald van; Larson, Mats; Badia, Santiago

DOI

[10.1007/s11831-023-09913-0](https://doi.org/10.1007/s11831-023-09913-0)

Publication date

2023

Document Version

Final published version

Published in

Archives of Computational Methods in Engineering

Citation (APA)

de Prenter, F., Verhoosel, C., Brummelen, H. V., Larson, M., & Badia, S. (2023). Stability and conditioning of immersed finite element methods: analysis and remedies. *Archives of Computational Methods in Engineering*, 30(6), 3617-3656. <https://doi.org/10.1007/s11831-023-09913-0>

Important note

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

Copyright

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.



Stability and Conditioning of Immersed Finite Element Methods: Analysis and Remedies

Frits de Prenter^{1,2} · Clemens V. Verhoosel³ · E. Harald van Brummelen³ · Mats G. Larson⁴ · Santiago Badia⁵

Received: 17 August 2022 / Accepted: 13 March 2023
© The Author(s) 2023

Abstract

This review paper discusses the developments in immersed or unfitted finite element methods over the past decade. The main focus is the analysis and the treatment of the adverse effects of small cut elements. We distinguish between adverse effects regarding the stability and adverse effects regarding the conditioning of the system, and we present an overview of the developed remedies. In particular, we provide a detailed explanation of Schwarz preconditioning, element aggregation, and the ghost penalty formulation. Furthermore, we outline the methodologies developed for quadrature and weak enforcement of Dirichlet conditions, and we discuss open questions and future research directions.

List of Symbols

| | |
|--|---|
| \mathcal{A} | Ambient or embedding domain |
| Ω | Problem (or physical) domain |
| Ω_h | Union of active elements |
| $\partial\Omega_D, \partial\Omega_N$ | Dirichlet and Neumann boundary |
| $n, \partial_n = n \cdot \nabla$ | Outer normal, normal derivative |
| T, T_Ω | Element, restriction of cut element T to Ω |
| h | Element size |
| h_{T_Ω} | Generalized thickness of cut element T |
| η_i | Volume fraction of T_i |
| η | Smallest volume fraction |
| η^* | Threshold volume fraction |
| $\mathcal{T}_h^A, \mathcal{T}_h$ | Ambient, active mesh |
| $\mathcal{T}_h^{\text{cut}}, \mathcal{T}_h^{\text{in}}, \mathcal{T}_h^{\text{ag}}$ | Cut, internal, aggregated mesh |
| F, \mathcal{F}_h | Element face, set of element faces |
| $\phi, \{\phi_i\}_{i=1}^N$ | Basis function, basis |
| p | Discretization order |

| | |
|--|--|
| $\mathcal{P}_p(T)$ | Polynomial space of order p over T |
| $\mathcal{Q}_p(T)$ | Tensor product space of order p over T |
| V_h | Approximation space |
| V_h^{ag} | Aggregated approximation space |
| N | System size |
| u | Analytical solution |
| u_h | Finite element solution |
| v_h, w_h | Finite dimensional functions |
| $\pi_h(\cdot)$ | Interpolation operator |
| $a_h(\cdot, \cdot)$ | Bilinear operator (2.7a) |
| $l_h(\cdot)$ | Linear operator (2.7b) |
| $s_h(\cdot, \cdot)$ | Ghost-penalty operator |
| $[\![\cdot]\!]$ | Jump operator (4.6) |
| $[\cdot]$ | Difference operator (4.15) |
| β | Nitsche parameter |
| τ | Ghost-penalty parameter |
| $\ \cdot \ _\Omega$ | L^2 norm over domain Ω |
| $(\cdot, \cdot)_\Omega$ | L^2 inner product over domain Ω |
| $\ \cdot \ _{a_h}$ | $a_h(\cdot, \cdot)$ operator norm |
| $\ \cdot \ _\beta$ | β -norm (named energy norm in [1]) |
| $\ \cdot \ _h, \ \cdot \ _{h,\star}$ | Unstabilized, stabilized energy norm |
| $\ell^2(N)$ | Vector space of size N |
| $\mathbf{x}, \mathbf{y}, \mathbf{z}$ | Coefficient vectors of u_h, v_h, w_h |
| \mathbf{A} | System matrix |
| \mathbf{b} | Right-hand-side vector |
| \mathbf{B} | Preconditioner |
| $\ \cdot \ _2$ | Euclidean vector or matrix norm |
| $\ \cdot \ _{\mathbf{A}}$ | \mathbf{A} -matrix vector energy norm |
| $\kappa(\cdot)$ | Condition number |

✉ Frits de Prenter
f.deprenter@tudelft.nl

¹ Faculty of Aerospace Engineering, Delft University of Technology, 2600 AA Delft, The Netherlands

² Reden - Research Development Netherlands, 7555 RJ, Hengelo (Ov.), The Netherlands

³ Department of Mechanical Engineering, Eindhoven University of Technology, 5600 MB Eindhoven, The Netherlands

⁴ Department of Mathematics and Mathematical Statistics, Umeå University, 901 87 Umeå, Sweden

⁵ School of Mathematics, Monash University, Clayton, Victoria 3800, Australia

1 Introduction


Over the past decades, the finite element method (FEM) has become an essential tool in scientific research and engineering. In its standard form, the finite element method requires the construction of a mesh that fits to the boundary of the considered geometry. For problems of practical interest, such a boundary-fitting mesh is constructed using mesh generators. Automatically generating boundary-fitting meshes can lack robustness, however, in the sense that manual intervention is required to, for example, repair distorted elements or non-matching surfaces. This is particularly the case when the geometry is very complex, such as for (possibly non-water-tight) CAD objects with many patches and trimming curves, geometries presented in the form of scan data, or settings in which frequent remeshing is required (for example in fluid–structure interactions). For such problems, immersed methods have been demonstrated to be capable of establishing a more efficient analysis pipeline, as illustrated by the examples in Fig. 1.

The pivotal idea of immersed finite element methods is to embed a complex geometry into a geometrically simple ambient domain, on which a regular mesh can be built easily. Basis functions defined on this ambient-domain mesh are then restricted to the problem geometry, after which the solution is approximated by a linear combination of these restricted basis functions. Although this discretization is conceptually straightforward, non-standard treatment of various aspects (which are discussed below) is required on account of the fact that the mesh does not fit the boundaries of the geometry. Using a wide variety of techniques to treat these non-standard aspects, the concept of immersed finite element methods has successfully been applied in a broad range of fields, such as solid mechanics [2–10]; shell analysis [9–15]; interface problems [16–25]; fluid mechanics [26–34]; fluid–structure interaction [35–43], in particular fluid–structure interaction for biomedical applications [44–47]; scan-based analysis of both man-made and

biological materials [48–55]; shape and topology optimization [56–62]; and many more.

The immersed analysis concept was originally proposed in the context of the finite difference method by Peskin [65] in 1972. This immersed boundary method (IBM) and its enhancements have been employed in a wide range of applications ever since (see, e.g., Ref. [66] for a contemporary review). The application of the immersed element concept in the finite element setting can be traced back to the work on the partition of unity method [67, 68], generally referred to as the generalized or extended finite element method (GFEM [69, 70] or XFEM [71, 72]), where elements are cut in order to construct enrichment functions. The concept of cutting finite elements as an unfitted meshing technique was pioneered by Hansbo [73]. This work can be considered as the first instance of an immersed finite element method. The pace in the development and impact of immersed finite element methods increased significantly with the introduction of the finite cell method (FCM) [2, 9, 74, 75], which combines the cut element concept with higher-order basis functions, and CutFEM [20, 26, 76–78], which generally employs the ghost penalty to enhance numerical stability [76, 79]. Besides these prominent immersed FEM techniques, other notable examples are the aggregated finite element method (AgFEM) [80, 80, 81], the Cartesian grid finite element method (cgFEM) [57, 82], weighted extended B-splines (WEB-splines) [83, 84] and immersed B-splines (i-splines) [85]. Discontinuous Galerkin methods can readily be used on cut meshes after aggregation of cells, since they can be posed on polytopal meshes [86]. In recent years, the immersed analysis concept has been considered in conjunction with isogeometric analysis [87, 88] (often referred to as iga-FCM [9] or immersogeometric analysis [45]). In this setting, immersed methods have been demonstrated to be capable of leveraging the advantageous properties of the spline basis functions used in isogeometric analysis, while enhancing the versatility of the simulation

Fig. 1 Examples of applications of immersed FEM. **a** shows the stress in an aluminum die cast gearbox housing [63]. The geometry is implicitly defined from CT data of a product to investigate stress concentrations around pores. **b** depicts the stress in a specimen of trabecular bone, rendered from CT data [64]



workflow for cases where boundary-fitting spline geometries are not readily available, e.g., in scan-based analyses.

Immersed finite element methods are typically confronted by three computational challenges in comparison to standard mesh-fitting finite elements, *viz*: (i) the numerical evaluation of integrals over cut elements; (ii) the imposition of (essential) boundary conditions over the immersed or unfitted boundaries; and (iii) the stability of the formulation in relation to small cut elements. A myriad of advanced techniques has been developed over the past decades to resolve these challenges, which has made immersed finite element techniques a competitive simulation strategy for a wide range of problems. This review focuses on the third challenge, i.e., the effects of small cut elements on the performance of immersed finite element methods. We restrict ourselves to a high-level consideration of the first two of these challenges, and we refer the reader to, e.g., Ref. [75] for a detailed review on these topics.

Small cut elements typically give rise to stability and conditioning problems. In this article we review three prominent methods to resolve these problems, which have been developed in recent years, *viz*: (Schwarz) preconditioning, ghost-penalty stabilization (commonly used in CutFEM), and aggregation of cut elements (commonly used in AgFEM). In the discussion of these techniques, it is important to note that small cut elements do not only affect the conditioning of the problem, but also the accuracy of the solution. More specifically, direct application of Nitsche's method for the weak imposition of essential boundary conditions requires a Nitsche parameter that scales inversely proportional with the size (in particular the thickness) of the cut element. As observed in, e.g., [1], the unboundedness of this parameter can deteriorate the accuracy of the solution. When using preconditioning techniques to resolve the small-cut-element problem, it is important to realize that these techniques do not resolve this potential issue regarding the accuracy. In contrast to preconditioning techniques, ghost penalty stabilization and aggregation ensure well-posedness with a Nitsche parameter inversely proportional to the ambient-domain mesh size, in addition to controlling the condition number. This makes the immersed finite element approximation using these stabilization techniques robust with respect to the cut element configurations, and preserves the error estimates of boundary-fitted finite element methods. It should be noted that, besides ghost penalty stabilization and aggregation, several other techniques have been developed to resolve the stability problems on small cut elements. However, in general, these do not simultaneously resolve the conditioning problem. An overview of these techniques is also presented in this review.

This review article has three objectives, *viz*: (i) to clarify that the small-cut-element problem is multi-faceted; (ii) to present the different techniques in an accessible theoretical

framework, so that the implications of practical choices become apparent to a non-expert audience; and (iii) to provide a comparison of the techniques for conditioning and stabilization of immersed finite element methods. With the large number of immersed finite element techniques available, naturally comes the luxury problem of choosing which technique is most suitable in a particular situation. The rigorous theoretical underpinning of the methods as reviewed in this article is instrumental to aiding in the selection of a particular method, as it provides a fundamental understanding of the relation between small cut elements, conditioning, and stability and accuracy.

It should be noted that immersed methods are not the only techniques to create a robust workflow to deal with complex or implicitly defined geometries. One alternative is the shifted boundary method [89–92]. The shifted boundary method aims to replace an immersed problem with a similar boundary-fitted problem on the interior element mesh, by projecting boundary conditions from the real (unfitted) boundary to the interior element boundaries. This concept was already introduced in [93], and is also applied in [62]. This method bypasses the aforementioned three computational challenges of immersed finite element methods, but instead introduces other challenges, such as a non-trivial treatment of boundary conditions (including the projection of the boundary data), and a non-obvious geometrical treatment. Another approach is to use hybridizable techniques on unfitted meshes [94, 95]. Hybridizable methods can naturally be posed on polytopal meshes, giving additional geometrical flexibility compared to standard finite element methods. As a result, these methods can readily be used on the meshes obtained after the intersection of the boundary representation and background mesh and possibly after aggregation of elements. The impact of small cut elements and small cut faces (these schemes add unknowns on the mesh skeleton) on stability and condition numbers has only been studied very recently in the context of hybrid high-order methods, see [95]. A detailed discussion of these alternative techniques is beyond the scope of this work.

This article is structured as follows. In Sect. 2 we introduce the basic formulation based on a model problem and discuss the developments regarding the three computational challenges associated with immersed finite element methods. In Sect. 3 a compact analysis of the ill-conditioning problem is presented, and Schwarz preconditioning is discussed as a natural technique to resolve this problem. Section 4 then considers stabilization techniques, specifically the ghost-penalty method and aggregation technique, and presents a theoretical framework required to analyze the stability and conditioning properties of these techniques. A discussion on the current state of the field and concluding remarks are finally presented in Sect. 5. Note that the results presented in this manuscript are reproduced from previous publications

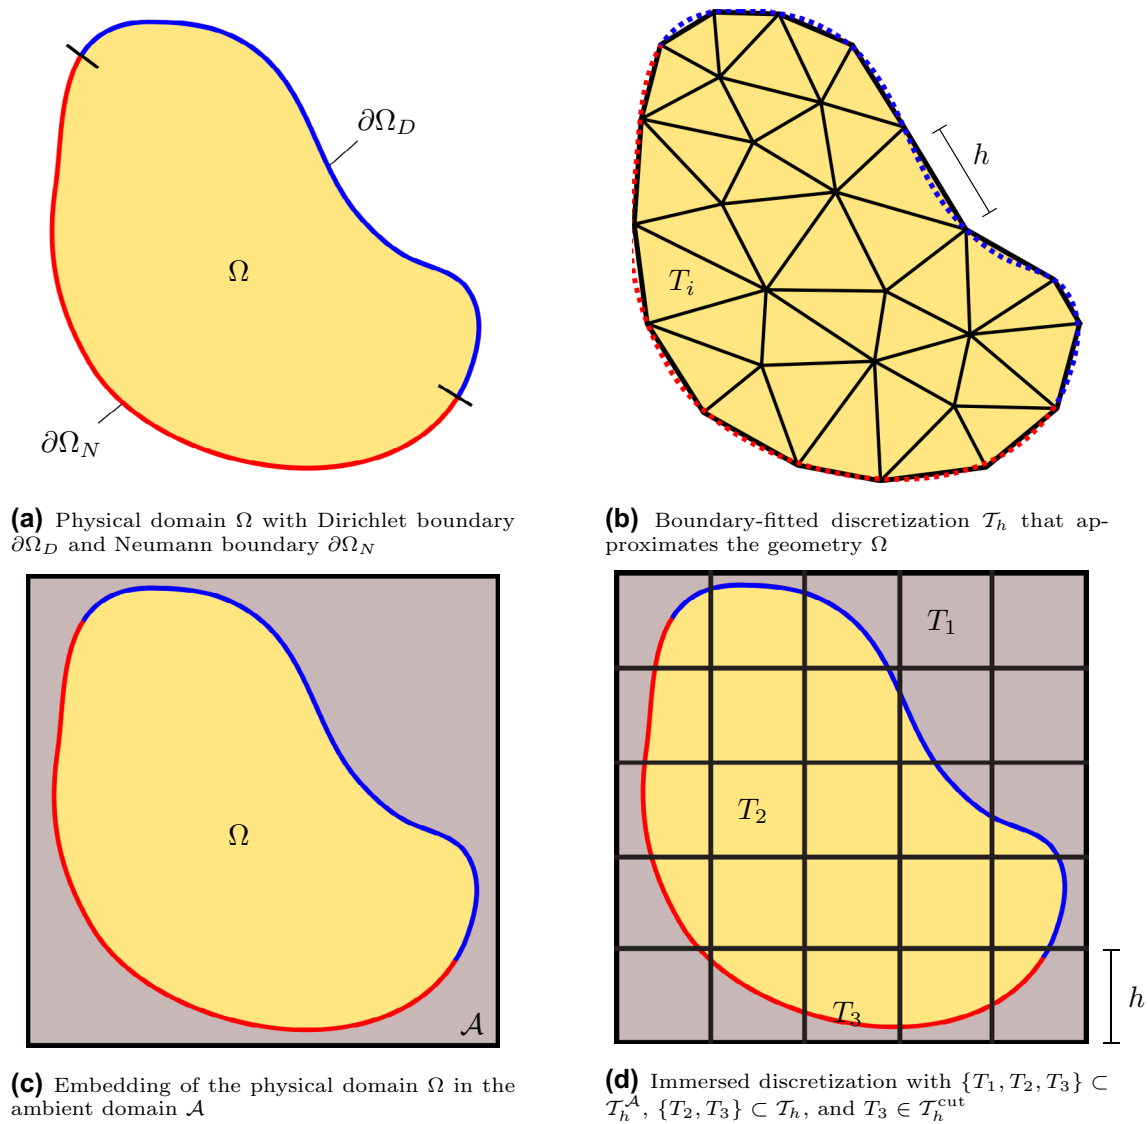


Fig. 2 Schematic representation the domain Ω , a boundary-fitted discretization, the embedding in \mathcal{A} , and an unfitted discretization

by the authors, which are referenced in the text or in the captions.

2 Immersed Finite Element Methods

In this section we introduce the immersed finite element framework. Section 2.1 presents the concept of immersed finite element methods and specifies the formulations based on a model problem. In Sect. 2.2 the most prominent challenges of immersed finite element methods, compared to standard boundary-fitted finite element methods, are discussed.

2.1 Formulations

As a model problem, we consider the Poisson equation on the domain $\Omega \subset \mathbb{R}^d$, with $d \in \{2, 3\}$ being the number of spatial dimensions (Fig. 2). The domain Ω has a boundary $\partial\Omega$ with outward-pointing unit normal vector n . The boundary consists of complementary parts $\partial\Omega_D$ and $\partial\Omega_N$ on which Dirichlet (or essential) and Neumann (or natural) conditions are prescribed with boundary data g_D and g_N , respectively. The field variable $u : \Omega \rightarrow \mathbb{R}$ is subject to the strong formulation

$$(\text{Strong}) \begin{cases} -\Delta u = 0 & \text{in } \Omega \\ u = g_D & \text{on } \partial\Omega_D \\ \partial_n u = g_N & \text{on } \partial\Omega_N \end{cases} \quad (2.1)$$

where $\partial_n = n \cdot \nabla$ denotes the normal gradient operator.

In boundary-fitted finite elements, the domain Ω is subdivided into elements T (generally at the expense of a geometrical error), which together comprise the mesh \mathcal{T}_h (Fig. 2b). The size of element T is denoted by h_T and a global mesh size parameter is defined as $h = \max_{T \in \mathcal{T}_h} h_T$. This manuscript only considers quasi-uniform discretizations, that is $h_T \approx h \forall T \in \mathcal{T}_h$. On the mesh, shape functions are introduced, which form the basis of the approximate solution u_h . Integration is performed via standard quadrature rules for polynomials on simplices or hypercubes. Homogeneous Dirichlet conditions are usually imposed strongly by removing the basis functions with support on $\partial\Omega_D$. The span of the remaining basis functions is referred to as $V_{h,0} \subset H_0^1(\Omega)$. The Bubnov-Galerkin finite element method imposes inhomogeneous boundary conditions through a so-called lifting function, $\ell(g_D) \in H^1(\Omega)$, which is generally constructed using the removed basis functions. The Bubnov-Galerkin finite element formulation corresponding to the strong form (2.1) can then be condensed into

$$(\text{Boundary-fitted FE}) \begin{cases} \text{find } u_h = w_h + \ell(g_D) \text{ for } w_h \in V_{h,0} \text{ such that:} \\ a(w_h, v_h) = l(v_h) \quad \forall v_h \in V_{h,0} \end{cases} \quad (2.2)$$

where the symmetric bilinear form $a(w_h, v_h) = \int_{\Omega} \nabla w_h \cdot \nabla v_h \, dV$ is continuous and coercive on $H_0^1(\Omega)$, and the linear form $l(v_h) = \int_{\partial\Omega_N} g_N v_h \, dS - \int_{\Omega} \nabla \ell(g_D) \cdot \nabla v_h \, dV$ is also continuous on $H_0^1(\Omega)$. These conditions are sufficient to guarantee well-posedness of the problem (2.2) to find $w_h \in H_0^1(\Omega)$ and $u_h \in H^1(\Omega)$ for sufficiently smooth boundary data [96, 97]. Because of the coercivity of the bilinear form, it induces the operator norm $\|v_h\|_a^2 = a(v_h, v_h)$, which is identical to the $H_0^1(\Omega)$ -seminorm for the boundary-fitted finite element formulation of the Poisson problem.

In immersed finite element methods, the domain Ω is embedded into an ambient domain \mathcal{A} (Fig. 2c). Instead of generating a boundary-fitted partitioning of the domain Ω , the ambient domain is partitioned by the background mesh $\mathcal{T}_h^{\mathcal{A}}$ (Fig. 2d). Since the ambient domain is geometrically simple, mesh generation is trivial. But, as the element boundaries do not coincide with the boundaries of the domain Ω , evaluating integrals requires dedicated procedures (see Sect. 2.2.1). Hence, the complexity of the geometry is essentially captured by the integration procedure, instead of by the mesh generation. This geometrical simplification also has implications on the Galerkin formulation.

To specify the immersed finite element formulation, we define the active mesh \mathcal{T}_h as the set of elements T which intersect the domain Ω , i.e.,

$$\mathcal{T}_h = \{T \in \mathcal{T}_h^{\mathcal{A}} \mid T \cap \Omega \neq \emptyset\} \quad (2.3)$$

where $T_{\Omega} = T \cap \Omega$. Furthermore, the active domain is defined as the union of all the active elements $\Omega_h = \cup_{T \in \mathcal{T}_h} T \supseteq \Omega$. The set of active elements that are cut by (and therefore intersect) the boundary is defined as

$$\mathcal{T}_h^{\text{cut}} = \{T \in \mathcal{T}_h \mid T \setminus T_{\Omega} \neq \emptyset\} \quad (2.4)$$

and the set of internal elements that are fully supported on the domain Ω is defined as

$$\mathcal{T}_h^{\text{in}} = \{T \in \mathcal{T}_h \mid T \subset \Omega\} \quad (2.5)$$

Analogous to the boundary-fitted case, shape functions are introduced on the active mesh \mathcal{T}_h and an approximate solution u_h is formed as a linear combination of the restriction of the shape functions to Ω . The space spanned by these functions is referred to as V_h . In contrast to the boundary-fitted finite element formulation, in the unfitted setting it is generally not feasible to create a finite dimensional subspace of $V_{h,0} \subset H_0^1(\Omega)$, which is needed to strongly impose Dirichlet conditions. Instead, Dirichlet conditions are generally imposed weakly. The most common approach for weakly enforcing Dirichlet boundary conditions is via Nitsche's method [98], which gives rise to the immersed finite element formulation

$$(\text{Immersed FE}) \begin{cases} \text{find } u_h \in V_h \text{ such that:} \\ a_h(u_h, v_h) = l_h(v_h) \quad \forall v_h \in V_h \end{cases} \quad (2.6)$$

where the bilinear and linear form are defined as

$$a_h(u_h, v_h) = \int_{\Omega} \nabla u_h \cdot \nabla v_h \, dV + \int_{\partial\Omega_D} (\beta u_h v_h - u_h \partial_n v_h - v_h \partial_n u_h) \, dS \quad (2.7a)$$

$$l_h(v_h) = \int_{\partial\Omega_N} g_N v_h \, dS + \int_{\partial\Omega_D} (\beta g_D v_h - g_D \partial_n v_h) \, dS \quad (2.7b)$$

The parameter β , commonly referred to as the Nitsche or penalty parameter, must be chosen large enough to ensure coercivity of the bilinear form $a_h(\cdot, \cdot)$ in the discrete space V_h . A sufficient condition is $\beta > \max_{v_h \in V_h} \|\partial_n v_h\|_{\partial\Omega_D}^2 / \|\nabla v_h\|_{\Omega}^2$, which can be computed by solving a generalized eigenvalue problem [99]. This can result in arbitrarily high values of β over the entire boundary $\partial\Omega_D$. It is therefore common to select a local, element-wise, Nitsche parameter satisfying $\beta|_T > \max_{v_h \in V_h|_T} \|\partial_n v_h\|_{T \cap \partial\Omega_D}^2 / \|\nabla v_h\|_{T_{\Omega}}^2$ by solving a local

generalized eigenvalue problem for each element. Based on dimensional considerations, one can infer that $\beta|_T$ should be inversely proportional to a generalized thickness, h_{T_Ω} , of T_Ω normal to $T \cap \partial\Omega_D$ [100, 101]. We refer to h_{T_Ω} as a *generalized* thickness, because T_Ω can be irregularly shaped and, consequently, a well-defined length scale cannot generally be provided. While, with an element-wise parameter, a single small cut element does not result in a high value of β over the entire boundary, an element-wise parameter is still locally unbounded, which can lead to numerical issues [1]. With appropriate stabilization (see Sect. 4), the properties of the immersed formulation revert to those of the boundary-fitted setting considered in the original paper by Nitsche [98], and a global parameter inversely proportional to the mesh width of the background mesh suffices. In the remainder of this manuscript, both element-wise and global Nitsche parameters will be indicated by β and the notation $a_h(\cdot, \cdot)$ and $l_h(\cdot)$ will be used in both stabilized and unstabilized formulations, with the choice of the Nitsche parameter following from the context. Similar to the boundary-fitted setting, the coercive and bounded bilinear form a_h induces an equivalent operator norm $\|v_h\|_{a_h}^2 = a_h(v_h, v_h)$ in the discrete space V_h . This operator norm coincides with the matrix energy norm of a corresponding coefficient vector, which renders it useful in the analysis of condition numbers of linear-algebraic systems emerging from immersed formulations.

Multiple variants of the immersed finite element formulation in (2.6)–(2.7) exist, the most prominent of which will be discussed in Sect. 2.2.2. A particularly noteworthy variation to (2.7) is the penalty method, which corresponds to the case where the terms that involve normal gradients in the above-mentioned operators are left out. This makes the method formally inconsistent with the strong form (2.1), but simplifies the implementation and ensures coercivity independent of the parameter β . In this review, we do not discuss the penalty method in detail and focus on the application of Nitsche's method. In general, the use of the penalty method instead of Nitsche's method has a negligible impact on the conditioning, but does affect the accuracy of the solution.

To solve the immersed finite element formulation (2.6) numerically, it is recast as a linear algebra problem

$$\mathbf{Ax} = \mathbf{b} \quad (2.8)$$

where the components of the matrix $\mathbf{A} \in \mathbb{R}^{N \times N}$ and right hand side vector $\mathbf{b} \in \mathbb{R}^N$ correspond to

$$\mathbf{A}_{ij} = a_h(\phi_j, \phi_i), \quad b_i = l_h(\phi_i) \quad (2.9)$$

with ϕ_i the i -th basis function and N the number of dimensions of the finite dimensional function space V_h . The approximate solution u_h is given by

$$u_h = \sum_i x_i \phi_i \quad (2.10)$$

with x_i the components of the coefficient vector $\mathbf{x} \in \ell^2(N)$ in equation (2.8). In the remainder we employ two norms for this coefficient vector, viz the ℓ^2 vector norm $\|\mathbf{x}\|_2^2 = \mathbf{x}^T \mathbf{x}$ and, as \mathbf{A} is Symmetric Positive Definite (SPD), the matrix energy norm

$$\|\mathbf{x}\|_{\mathbf{A}}^2 = \mathbf{x}^T \mathbf{A} \mathbf{x} \quad (2.11)$$

Note that, on account of (2.9), this matrix energy norm is equal to the operator norm, i.e., $\|\mathbf{x}\|_{\mathbf{A}} = \|u_h\|_{a_h}$.

In the remainder of this manuscript we focus our presentation on the single field Poisson problem (2.1), discretized by quasi-uniform meshes with C^0 -continuous piecewise polynomial basis functions of order p . Unless otherwise specified, results apply to both linear and higher-order discretizations. The presented analyses and methods naturally extend to vector-valued problems that can be expressed as a Cartesian product of scalar fields (i.e., one field per space dimension). Mixed formulations and discretizations with local refinements are not discussed in detail, but, if not stated otherwise, the provided insights also carry over *mutatis mutandis* to these cases. Particular aspects of interface problems and transient problems are discussed, respectively, in Remarks 2.2 and 2.3 in Sect. 2.2.3. With respect to conditioning, maximum continuity splines as commonly used in isogeometric analysis [87, 102] can behave differently from C^0 -continuous bases. Therefore, Sect. 3 also considers B-spline bases as a special case.

2.2 Challenges in Immersed Finite Elements

Although the immersed finite element formulation introduced above fits within the general framework of the conventional finite element method, the application of immersed finite elements involves the consideration of various specific challenges. In this section we discuss the most prominent of these, viz: (i) numerical evaluation of integrals over cut elements, (ii) imposition of Dirichlet boundary conditions, and (iii) stability and conditioning of the formulation.

2.2.1 Cut-Element Integration

In boundary-fitted FEM, integrals over the domain Ω are split into element-wise integrals. The elements generally correspond to polygons such as simplices or hypercubes, and the integrands are usually element-wise polynomials. For this reason, standard quadrature rules can readily be used. Integration in immersed methods is more involved, as the integration procedure should adequately approximate integrals on, in principle, arbitrarily shaped cut elements.

This problem closely relates to the special treatment of discontinuous integrands in enriched finite element methods such as XFEM and GFEM.

In immersed finite element methods the geometry representation is independent of the mesh. In general, there are two ways to represent the geometry, *viz* implicit representations (e.g., voxel data [21, 103] or a level set function [51]) and explicit boundary representations (e.g., spline surfaces [3], B-rep objects [9], or isogeometric analysis on trimmed CAD objects [8, 12]). For all geometry representations, dedicated techniques are required to evaluate volume integrals over the intersection of active elements with the domain. A myriad of such integration procedures has been developed over the years in the context of immersed FEM (see [75, 104] for reviews/comparisons) and enriched FEM (see [105]), an overview of which is presented below. Techniques to integrate over trimmed boundaries are strongly dependent on the geometry description. In explicit representations, the geometry description itself can sometimes be leveraged, whereas generally a boundary-reconstruction procedure is required in the case of implicit boundary representations. The reader is referred to, e.g., Refs. [75] for a discussion regarding the various techniques to handle different geometry representations, and the related problem of integrating over unfitted boundaries.

The cut element volume integration techniques can be categorized as:

- *Octree subdivision*: The general idea of octree (or quadtree in two dimensions) integration is to capture the geometry of a cut element by recursively bisecting sub-cells that intersect with the boundary of the domain, as illustrated in Fig. 3. At every level of this recursion, sub-cells that are completely inside the domain are preserved, while sub-cells that are completely outside of the domain are discarded. This cut element subdivision strategy was initially proposed in the context of the finite cell method (FCM) in [2] and is generally appraised for its simplicity and robustness with respect to cut element configurations. Octree integration has been widely adopted in immersed FEM, see, e.g., [45, 51, 75, 106, 107]. Various generalizations and improvements to the original octree procedure have been proposed, of which the consideration of tetrahedral cells [52, 108], the reconstruction of the unfitted boundary by tessellation at the lowest level of bisectioning [51], and the consideration of variable integration schemes for the sub-cells [104], are particularly noteworthy. Despite the various improvements to the original octree strategy, a downside of the technique remains the number of integration sub-cells (and consequently the number of integration points) that result from the procedure, especially in three dimensions




Fig. 3 Illustration of the octree procedure to integrate cut elements, in which integration sub-cells that intersect with the immersed boundary are recursively bisected

and with high-order bases, where the refinement depth needs to be increased under mesh refinement to reduce the integration error with the same rate as the approximation error [109].

- *Cut element reparameterization*: Accurate cut element integration schemes can be obtained by modifying the geometry parameterization of cut elements in such a way that the immersed boundary is fitted. This strategy was originally developed in the context of XFEM by decomposing cut elements into various sub-cells with only one curved side and then to alter the geometry mapping related to the curved sub-cell to obtain a higher-order accurate integration scheme [105]. This concept has been considered in the context of implicitly defined geometries (level sets) [110–113], the NURBS-enhanced finite element method (NEFEM) [114, 115], the Cartesian grid finite element method (cgFEM) [57, 82], and the mesh-transformation methodology presented in [116]. In the context of the finite cell method, the idea of cut element reparameterization has been adopted as part of the *smart octree integration* strategy [117–119], where a boundary-fitting procedure is employed at the lowest level of octree bisectioning in order to obtain higher-order integration schemes for cut elements with curved boundaries. Reparameterization procedures have the potential to yield accurate integration schemes at a significantly lower computational cost than octree procedures, but generally compromise in terms of robustness with respect to cut element configurations.

- *Polynomial integration*: Provided that one can accurately evaluate integrals over cut elements (for example using octree integration), it is possible to construct computationally efficient integration rules for specific classes of integrands. In the context of immersed finite element methods, it is of particular interest to derive efficient cut-element integration rules for polynomial functions. The two most prominent methods to integrate polynomial functions over cut elements are *moment fitting techniques* [119–126], in which integration point weights and (possibly) positions are determined in order to yield exact quadrature rules, and *equivalent polynomial methods* [127, 128], in which a non-polynomial (e.g., discontinuous) integrand is represented by an equivalent polynomial which can then be treated using standard integration procedures. Such methods have been demonstrated to yield efficient quadrature rules for a range of scenarios. A downside of such techniques is the need for the evaluation of the exact integrals (using an adequate cut-element integration procedure) in order to determine the optimized integration rules. This can make the construction of such quadrature rules computationally expensive, which makes them more suitable in the context of time-dependent and non-linear problems (with fixed boundaries and interfaces), for which the construction of the integration rule is only considered as a pre-processing operation (for each cut element) and the optimized integration rule can then be used throughout the simulation. Another downside is that some of these methods can result in negative quadrature weights, which can cause instabilities as discussed in [126]. In the group of Dominik Schillinger at TU Darmstadt, work is currently being done to use neural networks for the computation of integration point weights to accelerate this process.
- *Dimension-reduction of integrals*: Depending on the problem under consideration, it can be possible to reformulate volumetric integrals over cut elements by equivalent lower-dimensional integrals. This approach is advantageous from a computational effort point of view, as the equivalent integrals are generally less costly to evaluate. A reformulation of volume integrals in terms of boundary integrals has been proposed in the context of XFEM in [129] and in the immersed FEM setting in [130]. Dimension-reduction approaches for high-order quadrature with implicitly defined surfaces are presented for hexahedral and tetrahedral background elements in, respectively, [131] and [132]. Both these techniques rely on a reduction of integrals to one-dimensional integrations and provide strictly positive weights. The methodology proposed in [133] (and references therein) provides

closed-form formulas for the integration of monomials on convex and nonconvex polyhedra and the extension to curved domains, relying on a reduction of integrals up to vertex evaluations. A downside of dimension-reduction techniques is that they are less general than standard quadrature rules.

- *Parameter optimization*: Various strategies have been proposed to optimize the parameters of the cut element integration techniques discussed above, most notably for octree subdivision techniques. In [104, 109] algorithms are proposed to select the integration order on the different levels of sub-cells. Ref. [126] presents a methodology to reduce the number of integration points in a manner similar to moment fitting techniques. These optimization techniques have demonstrated that reducing the number of integration points does not necessarily compromise the accuracy of the simulations. This is theoretically supported by Strang's first lemma [96, 134, 135], which indicates that integration does not need to be exact in order to attain (optimal) convergence [134]¹. It should be mentioned that this lemma is also considered in the context of CutFEM in e.g., [109, , 137].

In the selection of an appropriate cut element integration scheme one balances robustness (with respect to cut element configurations), accuracy, and expense. If one requires a method that automatically treats a wide range of cut element configurations and is willing to pay the price in terms of accuracy and computational effort, octree integration is the compelling option. For moderate accuracy, quadratures that rely on exact monomial integrations and only involve vertex evaluations¹³⁶ are appealing in terms of accuracy and robustness. If the accuracy and computational-expense requirements are more stringent and the range of configurations is suitably restricted, alternative techniques such as cut element reparameterization are attractive. In the case of implicit boundary representations, an additional consideration in the selection of the cut element integration scheme is whether or not it is required to obtain a parameterization of (or integration scheme on) the unfitted boundary. Parameter optimization procedures can aid in fine-tuning the balance between robustness, accuracy and expense. The appropriateness of the various techniques also depends on the way in which the geometry is represented (implicit vs explicit), as this can have a substantial impact on the implementation of a particular technique.

¹ As formulated by Strang and Fix in 1973 [134]: "What degree of accuracy in the integration formula is required for convergence? It is not required that every polynomial which appears be integrated exactly."

2.2.2 Dirichlet Boundary Condition Imposition

Since the (immersed) boundary of the physical domain does not coincide with the background mesh, the imposition of boundary conditions in immersed methods requires special consideration. Given that a parameterized boundary representation for integration over the boundary exists, Neumann (or natural) boundary conditions can be imposed weakly, in the same way as done in boundary-fitted FEM. The imposition of Dirichlet (or essential) boundary conditions in immersed FEM is not as straightforward, however. Because of the disparity between the background grid and the physical domain, basis functions defined on the background mesh are not generally interpolatory on the unfitted boundary. This precludes strong imposition of Dirichlet conditions as in boundary-fitted FEM. Therefore, boundary conditions in immersed FEM are usually imposed weakly. Different techniques for the imposition of Dirichlet boundary conditions on unfitted boundaries exist, the most prominent of which are:

- *Penalty method:* The penalty method supplements the weak form of boundary-fitted FEM with a penalty term that penalizes differences between the approximate solution and the prescribed Dirichlet data. This approach, which has been applied in the pioneering work on the finite cell method [2], is generally considered as the most straightforward technique to impose Dirichlet conditions on unfitted boundaries. The formulation omits the boundary terms that arise from the partial integration in the derivation of the weak form—in boundary-fitted FEM these terms drop out as the test functions vanish on Dirichlet boundaries—such that a modeling error is introduced which yields an inconsistent formulation. For appropriately selected penalty parameters the inconsistency can be acceptable [138], making the penalty method effective for a broad class of immersed problems with complex geometries, see, e.g., [21, 103]. Nevertheless, the choice of an appropriate penalty parameter is challenging. A too small value does not adequately enforce the prescribed boundary conditions, while a too large value exacerbates the conditioning problems [8] and can lead to large, nonphysical, gradients on cut elements [1, 139, 140].
- *Nitsche’s method:* Nitsche’s method [98] can be considered as the consistent equivalent of the penalty method, as it retains the boundary gradients (sometimes referred to as the flux terms) in the weak formulation. Through appropriate scaling of the Nitsche (or penalty) parameter, a stable formulation is obtained, see, e.g., [141]. Nitsche’s method is a widely used technique for the weak imposition of boundary conditions in immersed finite element methods. An elegant aspect of Nitsche’s method is that the parameters can be computed per element [99], avoiding potential difficulties in the selection of a single global Nitsche parameter, see, e.g., [10]. The value of the Nitsche parameter should be inversely proportional to the thickness of the element,² and can become arbitrarily large for small cut elements [1]. This problem can be remedied by means of additional stabilization terms such as the ghost penalty, or by the aggregation of basis functions; see Sect. 4. Also, a nonsymmetric Nitsche method can be applied to avoid the need for stabilization [142–144], although this does affect the linear solver. Additionally, nonsymmetric Nitsche methods are not adjoint consistent, which means that these formulations result in suboptimal approximation properties in the $L^2(\Omega)$ -norm [142]. Some other variations of Nitsche’s method are presented in [145].
- *Lagrange multiplier techniques:* Dirichlet conditions on immersed boundaries can be enforced by supplementing the weak formulation with additional constraint terms, e.g., [141, 146]. In contrast to the penalty method and Nitsche’s method, in Lagrange multiplier techniques these constraint terms are associated with an auxiliary field variable which is defined over the unfitted boundary. This auxiliary field is referred to as the Lagrange multiplier field. Lagrange multiplier techniques result in a saddle point problem, which implies that the discrete Lagrange multiplier field needs to be selected in such a way that a stable system is obtained [147]. Examples of Lagrange multiplier type techniques for immersed FEM are presented in [39, 148–151]. While an advantage of Lagrange multiplier techniques over Nitsche’s method and the penalty method is that these do not require the selection of a parameter, the downsides are that additional degrees of freedom are introduced through the Lagrange multiplier field, and that a (inf-sup) stable discretization of the Lagrange multiplier field is generally non-trivial. Additionally, for many problems, the introduction of Lagrange multipliers changes the nature of the linear system from positive (semi-)definite to indefinite and breaks the diagonal dominance. This affects the applicability of iterative solvers (in particular this precludes the conjugate gradient method) and the factorization in sparse direct solvers.
- *Basis function redefinition:* An alternative class of techniques to impose Dirichlet conditions on immersed boundaries is based on the idea to redefine the basis functions in such a way that the modified (non-vanishing) basis functions are interpolatory on the unfitted bound-

² This holds for second-order elliptic problems like the Poisson problem (2.1). With higher-order equations some parameters require higher scaling rates. [99]

ary. This enables traditional strong imposition of boundary conditions, as is standard in boundary-fitted finite element methods. Prominent examples of methods based on this concept are WEB-splines [83, 84] and, more recently, i-splines [85]. The main advantage of these techniques is that they do not require modifications to the weak formulation in comparison to boundary-fitted FEM, but the algorithms to redefine the basis functions constitute an additional non-trivial component in the implementation and analysis.

In the selection of an appropriate technique for imposing Dirichlet conditions in immersed finite elements, consistency and accuracy requirements are a prominent consideration. If there are no stringent accuracy requirements, the penalty method is an attractive option on account of its simplicity. If this method does not meet the accuracy requirements, one can resort to consistent weak formulations, where particularly Nitsche's method strikes a suitable balance between accuracy and ease of implementation. Basis function redefinition strategies are an attractive alternative when there are reasons to enforce Dirichlet conditions in a strong manner, like in boundary-fitted finite element formulations.

2.2.3 Stability and Conditioning

Immersed discretizations that make use of finite element spaces defined on a background mesh generally suffer from the so-called *small-cut-element problem*. Conventional boundary-fitted finite element methods impose conditions on the shape and the size of the elements in the considered (family of) meshes. Such conditions can be (more-or-less) directly managed by the mesh-generation algorithm. In immersed finite element methods, on the other hand, one has no control over the shape and size of cut elements. Consequently, cut elements can have arbitrarily small intersections with the physical domain, which is commonly referred to as the small-cut-element problem. This loss of control in the immersed or unfitted setting can have extensive implications on the well-posedness and conditioning of the resulting discrete problem, unless the immersed method is judiciously formulated.

The challenge of stability and conditioning in immersed finite elements with respect to small cut elements is the main focus of this review. Frequently, this small-cut-element problem is considered as a single-faceted problem. In our opinion, however, there are two distinct (albeit strongly related) facets to the small-cut-element problem, *viz* stability and conditioning:

- *Stability* is related to the fact that the numerical formulation itself can be ill-posed in the immersed setting. The most clear example of this is the imposition of Nitsche's

method with standard unfitted finite element spaces. The Nitsche parameter required for coercivity tends to infinity under mesh refinement, which can lead to unbounded gradients on immersed or unfitted boundaries. It should be mentioned that the stability of immersed finite elements is closely related to the way in which essential boundary conditions are enforced, and, in the case of Nitsche's method, to the value of the Nitsche parameter.

- *Conditioning* is related to the linear algebraic problem of obtaining the solution of the discrete system that arises from an immersed finite element formulation. Even if the problem is properly defined from a stability perspective, e.g., with only Neumann conditions on unfitted boundaries, the resulting system can be arbitrarily ill-conditioned, which impedes the solution of the system. This is caused by functions that are only supported on small cut elements, for which the operator norm (that is equal to the matrix energy norm of the corresponding coefficient vector) is affected by the cut-element size, while the norm of the coefficient vector itself is not. This implies that the eigenvalues of the system matrix can be arbitrarily close to zero, depending on the cut-element configuration.

Several techniques have been developed to counteract the problems associated with small cut elements, encompassing treatments for both stability and conditioning issues. The effects of cut elements on the conditioning of the linear system are discussed in detail in Sect. 3.1. Section 3.2 provides an overview of tailored preconditioners that address this issue. Specific attention in this section is devoted to Schwarz preconditioners, that form a natural resolution to the conditioning problem. The stability of immersed discretizations is treated in detail in Sect. 4, which considers both the stability problem itself in Sect. 4.1 and the methodologies that have been developed to resolve it in Sect. 4.2. Two particular approaches, *viz* element aggregation and the ghost-penalty method, resolve the stability problems in a manner that yields enhanced coercivity in $H^1(\Omega_h)$ (i.e., on the union of all active elements) instead of just in $H^1(\Omega)$. Consequently, these approaches do not only guarantee stability, but simultaneously preclude conditioning problems. Element aggregation and the ghost-penalty method are discussed in detail in Sects. 4.3 and 4.4, respectively, and Sect. 4.5 presents a unified mathematical approach to analyze these techniques.

Remark 2.1 Interpretation from the perspective of norm equivalences In terms of the problem definitions presented above, stability and conditioning can be distinguished as different norm equivalences in the discrete space. Stability pertains to the strength of the norm equivalence between the $H^1(\Omega)$ -norm, which is a common measure for establishing the quality of a solution, and the operator norm, $\|\cdot\|_{a_h}$ (or

the equivalent β -norm or energy norm, which will be defined in Sect. 4). Conditioning, on the other hand, pertains to the strength of the norm equivalence of the matrix energy norm, $\|\cdot\|_A$ (which is equal to the operator norm of the corresponding discrete function), and the ℓ^2 vector norm, $\|\cdot\|_2$. In an unfitted discretization that contains small cut elements, the equivalences between these norms can become very weak, which indicates the potential of difficulties with respect to stability and/or conditioning. To ensure coercivity of the weak formulation when Nitsche's method is employed in an immersed setting without a dedicated stabilization technique, a (locally) very large Nitsche parameter is required. Such a very large Nitsche parameter, however, also causes the operator norm to be only very weakly bounded by the $H^1(\Omega)$ -norm [1]. Similarly, functions that are only supported on small cut elements have a small operator norm, while the size of the cut elements does not affect the ℓ^2 -norm of the coefficient vector. For this reason, the bound on the ℓ^2 -norm of the coefficient vector in terms of the matrix energy norm can be arbitrarily weak [152].

Remark 2.2 Interface problems Immersed formulations of interface problems such as [16] form a special class when considering the stability of cut elements. In such problems, an interface that is not aligned with the mesh separates parts of the domain, instigating a jump in the (material) properties (e.g., stiffness or conductivity) and thereby in the coefficients of the equation. Transmission conditions at the interface generally specify continuity of the solution and the fluxes, resulting in a jump in the gradients over the interface. The solution on different sides of the interface is therefore approximated by separate function spaces (defined on the same mesh covering the entire domain), with parts of basis functions that are intersected by the interface separately involved on both sides of the interface. This constructs a special case of immersed methods, because cut elements always have significantly large supports on either side of the interface. This can be leveraged with a weighted average flux in the bilinear formulation as presented in [18, 19, 73], resulting in a stable formulation with a bounded Nitsche parameter without the requirement for additional stabilization. Note that this formulation does not repair the conditioning problems and that the weighting of the fluxes is not effective for high-contrast problems. Besides formulations with weighted average fluxes, stable immersed interface problems can also be obtained by the techniques presented in Sect. 4, which additionally repair the conditioning. Particular examples of stable formulations with aggregation and ghost penalty are presented in [23, 24, 153] and [20, 25], respectively.

Remark 2.3 Stable explicit time integration While not considered in detail in this review, a related challenge is the

stability of explicit time integrators on unfitted grids containing small cut elements. Analogous to the system matrix, the eigenvalues of the (consistent) mass matrix cannot be bounded from below in immersed formulations, such that the stable time step can become arbitrarily small. Similar to the stability of the solution on small cut elements discussed above, this can generally be resolved by the function space manipulations discussed in Sect. 4.3 or by the addition of the weak stabilization terms as discussed in Sect. 4.4 to the mass matrix, see, e.g., [154]. Regarding the stability of explicit time integrators, also the investigations into the spectral behavior of Nitsche's method in both boundary-fitted and immersed settings presented in [155–157] are of particular interest. It should be mentioned that systems with lumped mass matrices and smooth (isogeometric) discretizations form an exception to the dependence of the stable time-step size on cut elements. For such systems, the eigenvalues of the lumped mass matrix on small cut elements scale more favorably with the size of small cut elements than the eigenvalues of the stiffness matrix. Therefore, stable explicit time integration can be performed without additional stabilization and with time steps dependent on the background element size, see [158, 159] for details.

3 Ill-Conditioning and Preconditioning

An integral aspect of finite element methods is to find the solution of the corresponding linear system of equations. For small systems this is generally done by a direct solver, which factorizes the linear system and directly computes the solution up to machine precision. The computational cost of direct solvers scales poorly with the size of the system, however, resulting in computation time and memory requirements that become prohibitive for large systems [160]. For this reason, large systems are generally solved by iterative solvers, the computational cost of which generally scales better with the size of the system [161]. The convergence of iterative solvers is strongly correlated with the conditioning of the system. Without tailored stabilization or preconditioning, systems derived from immersed finite element formulations are generally severely ill-conditioned, impeding the application of iterative solvers [152]. In Sect. 3.1, an analysis of the causes of ill-conditioning in immersed finite element systems is presented. Preconditioners which effectively resolve these conditioning problems are discussed in Sect. 3.2. Section 3.3 provides a detailed discussion about Schwarz preconditioners, which have commonly been applied to immersed finite element systems in recent years.

3.1 Conditioning Analysis

3.1.1 Condition Number

The conditioning of a linear system is often used as an indication of the complexity of solving that linear system by means of an iterative solver. An important property of a system matrix to indicate the conditioning is the (Euclidean) condition number, $\kappa(\mathbf{A})$, which is defined as the product of the norm of the system matrix and the norm of its inverse

$$\kappa(\mathbf{A}) = \|\mathbf{A}\|_2 \|\mathbf{A}^{-1}\|_2 \geq 1 \quad (3.1)$$

While the condition number is formally defined as the sensitivity of the solution to perturbations in the right hand side, it also relates to the convergence of iterative solvers. It is noted that the convergence of these solvers is dependent on more factors than simply the condition number, such as, e.g., the grouping of eigenvalues and the orthogonality of eigenvectors [161, 162].

Based on the equality between the matrix energy norm $\|\mathbf{y}\|_{\mathbf{A}}$ of a vector \mathbf{y} in the vector space $\ell^2(N)$ and the operator norm $\|v_h\|_{a_h}$ of the corresponding function $v_h = \sum_i y_i \phi_i$ in the isomorphic function space V_h , the norm of the (symmetric) system matrix \mathbf{A} can be written as

$$\begin{aligned} \|\mathbf{A}\|_2 &= \max_{\mathbf{y}} \frac{\|\mathbf{A}\mathbf{y}\|_2}{\|\mathbf{y}\|_2} = \max_{\mathbf{z}, \mathbf{y}} \frac{a_h(w_h, v_h)}{\|\mathbf{z}\|_2 \|\mathbf{y}\|_2} \\ &= \max_{\mathbf{y}} \frac{a_h(v_h, v_h)}{\|\mathbf{y}\|_2^2} = \max_{\mathbf{y}} \frac{\|v_h\|_{a_h}^2}{\|\mathbf{y}\|_2^2} \end{aligned} \quad (3.2)$$

with $w_h = \sum_i z_i \phi_i \in V_h$ and $\mathbf{z} \in \ell^2(N)$. Similarly, for the norm of the inverse of \mathbf{A} it can be written that

$$\|\mathbf{A}^{-1}\|_2 = \max_{\mathbf{y}} \frac{\|\mathbf{y}\|_2^2}{\|v_h\|_{a_h}^2} \quad (3.3)$$

The identities in (3.2) and (3.3) convey that the condition number of a linear system emanating from an (immersed) finite element formulation is determined by the tightness of the equivalence between, on one side, the vector norm $\|\mathbf{y}\|_2$ of the coefficient vector, and, on the other side, the operator norm $\|v_h\|_{a_h}$ of the corresponding function, on the isomorphic spaces $\ell^2(N)$ and V_h . The weakness of this norm equivalence for systems that contain small cut elements is the root cause of ill-conditioning in immersed finite element methods.

The norms defined above can be employed to provide an estimate, or lower bound, for the condition number of a system with a small cut element. To derive this estimate, we consider a two-dimensional cut-element scenario as indicated in Fig. 4. First, we define the volume fraction η_i of the cut element $T_i \in \mathcal{T}_h^{\text{cut}}$ as the ratio between the volume of




Fig. 4 Exemplary cut element geometry to illustrate the derivation of the condition number estimate. The vertex with $\xi = (0, 0)$ is indicated in blue

the intersection of the element with the domain Ω and the volume of the full element in the background grid

$$\eta_i = \frac{|T_i \cap \Omega|}{|T_i|} = \frac{|T_{i,\Omega}|}{|T_i|} \quad (3.4)$$

Second, we note that in a polynomial basis of order p it is generally possible to construct a function $v_h \in V_h$

$$v_h = \begin{cases} \frac{\xi_1^p \xi_2^p}{h^{2p}} & x \in T_i \\ 0 & x \notin T_i \end{cases} \quad (3.5)$$

with $\xi = (\xi_1, \xi_2)$ a local coordinate that has its origin $\xi = (0, 0)$ at the (not-cut) vertex of T_i , as indicated in Fig. 4, and where the corresponding coefficient vector \mathbf{y} in the vector space $\ell^2(N)$ has a norm that is approximately $\|\mathbf{y}\|_2 \approx 1$ depending on the employed basis (e.g., B-splines, Lagrange, or integrated Legendre). Under the assumption that element T_i is cut by a natural boundary such that the weak form does not contain boundary terms, the operator norm of v_h is given by

$$\begin{aligned} \|v_h\|_{a_h}^2 &= \frac{1}{h^{4p}} \int_{\xi_1=0}^{h\eta_i^{\frac{1}{2}}} \int_{\xi_2=0}^{h\eta_i^{\frac{1}{2}}} \left((p\xi_1^{p-1}\xi_2^p)^2 + (p\xi_1^p\xi_2^{p-1})^2 \right) d\xi_1 d\xi_2 \\ &= \frac{2p^2}{(2p-1)(2p+1)} \eta_i^{2p} \end{aligned} \quad (3.6)$$

This shows that while the ℓ^2 vector norm $\|\mathbf{y}\|_2$ is not affected by the volume fraction η_i , the operator norm $\|v_h\|_{a_h}$ and, by equivalence, the matrix energy norm scales with the volume fraction η_i . Therefore it follows that

$$\|\mathbf{A}^{-1}\|_2 = \max_{\mathbf{y}} \frac{\|\mathbf{y}\|_2^2}{\|\mathbf{y}\|_{\mathbf{A}}^2} = \max_{\mathbf{y}} \frac{\|\mathbf{y}\|_2^2}{\|v_h\|_{a_h}^2} \gtrsim \frac{1}{\eta_i^{2p}} \quad (3.7)$$

with \gtrsim denoting an inequality involving a constant (i.e., $a \gtrsim b$ indicates $a > Cb$ for some constant C). As the norm $\|\mathbf{A}\|_2$ does not depend on the volume fraction η_i , this results in a condition number $\kappa(\mathbf{A}) \gtrsim \eta_i^{-2p}$. In [152] this derivation is performed in an abstract setting for more general cut scenarios (under certain shape assumptions), for different numbers of dimensions, and with different boundary conditions. This results in the estimate of the condition number of linear systems derived from immersed finite element formulations of second order problems

$$\kappa(\mathbf{A}) \gtrsim \eta^{-(2p+1-2/d)} \quad (3.8)$$

with η denoting the smallest volume fraction in the system, i.e., $\eta = \min_i \eta_i$. The scaling relation (3.8) is numerically verified in [152] for different grid sizes, discretization orders, and for both C^0 -continuous and maximum continuity spline bases.

It can be noticed that the shape of the cut element is not included in the condition number estimate. That is because within certain shape-regularity assumptions (see [152] for details), the shape of the cut element is not a dominant factor in the order of magnitude of the smallest eigenvalue of the system matrix. When the system matrix is (diagonally) scaled, as will be discussed below, the shape of the cut element does play a role. It should be mentioned that, under the same shape-regularity assumptions, the derivation in [152] also establishes that when a local, element-wise, Nitsche parameter is employed, the condition number estimate is not affected by the type of boundary condition imposed on the (cut) boundary of the smallest cut element. When a global Nitsche parameter is employed, the value of this parameter is determined by the smallest cut element. While this does not affect the smallest eigenvalue of the system, this makes the largest eigenvalue dependent on the smallest cut element as well, resulting in even larger condition numbers.

3.1.2 Effect of Diagonal Scaling, Smoothness and the Cut-Element Shape

The previous paragraph considered the condition number of the system matrix \mathbf{A} as is, without any treatment. It is noted, however, that linear algebra solvers typically perform basic rescaling procedures, most notably diagonal

scaling operations such as Jacobi, Gauss-Seidel or SSOR preconditioning [163]. Consequently, it is important to not only consider the condition number independently, but also to have a closer look at the coefficient vector \mathbf{y} corresponding to a function with a very small operator norm $\|v_h\|_{a_h}$. As indicated in [152], there exist two mechanisms by which the operator norm of a basis function can be much smaller than the norm of the corresponding coefficient vector. First, it is possible that a basis function in itself is very small. In this case, simply the unit vector corresponding to that specific basis function will already cause a large condition number through (3.3) (i.e., with this unit vector taken as \mathbf{y} , the quotient in this equation will be very large). Second, on small cut elements it is possible that the dependence of certain basis functions on a specific parametric coordinate or on higher-order terms diminishes. This essentially reduces the dimension of the space spanned by the basis functions on the small cut element, such that these basis functions become almost linearly dependent (see Fig. 5). In this case, the small function v_h as in (3.5) cannot be represented by a unit vector, and the coefficient vector \mathbf{y} corresponding to the small function requires multiple nonzero entries. This is an important nuance in relation to the scaling of the system, as small eigenvalues that correspond to an (almost) unit vector are resolved by diagonal scaling, while small eigenvalues caused by almost linear dependencies are not.

Because of the above-discussed distinction between the classes of small eigenvalues in immersed finite element systems, there is a significant difference in conditioning between, on one hand, higher-order ($p \geq 2$) C^0 -polynomials, and, on the other hand, maximal continuity splines and linears (note that linears are a subclass of maximal continuity splines). This is because in the former case, almost linear dependencies will be formed on all small cut elements. For such systems, Jacobi preconditioning lowers the condition number, but the resulting systems are still ill-conditioned and the condition number still shows a dependence on the smallest cut element. For maximum continuity splines, the shape of small cut elements starts to play a role. On elements in which a vertex is contained in T_Ω , as indicated by element $T_{1,\Omega}$ in Fig. 5, a discretization with maximum continuity splines will only contain a single basis function of which the support is restricted to the small cut element. Therefore, a Jacobi preconditioner suffices to repair the conditioning with regard to that element. On cut elements where T_Ω does not contain a vertex, as indicated by element $T_{2,\Omega}$ in Fig. 5, the diminished dependence on the (in this case horizontal) parametric coordinate causes almost linear dependencies. Consequently, such elements still cause ill-conditioning even after diagonal rescaling of immersed systems with maximum continuity splines. Because of this dependence on the continuity of




Fig. 5 Illustration of an unfitted boundary with different types of cut elements. Element $T_{1,\Omega}$ contains a vertex. With maximal continuity splines, only one basis function is supported on this element only (for a linear basis, the node corresponding to this function is indicated by the blue dot). Rescaling the basis functions will resolve the conditioning problems for this type of cut element. Element $T_{2,\Omega}$ does not contain a vertex, and multiple basis functions are only supported on this element (with a linear basis, these are the basis functions corresponding to the nodes indicated with the green dots). If the volume fraction of this element is very small, the dependence of the basis functions on the horizontal coordinate (relative to the dependence on the vertical coordinate) diminishes, such that the basis functions describe essentially the same degree of freedom and become linearly dependent. Therefore, rescaling the basis functions will not resolve the conditioning problems on this type of cut element. In a similar manner, the relative contribution of higher-order terms diminishes relative to the contribution of linear terms on small cut elements. As a result, with higher-order discretizations that are not of maximal continuity, almost linear dependencies generally occur on small cut elements of any shape, while these only occur on specifically cut elements with discretizations of maximal continuity

the basis, a scaled linear system derived from an unfitted discretization with maximum continuity splines will contain far fewer problematically small eigenvalues than a scaled linear system derived from an unfitted discretization with C^0 -continuous basis functions of the same degree $p \geq 2$. Depending on the size of the system and the smoothness of the boundary, the number of problematic eigenvalues in an immersed isogeometric discretization in conjunction with diagonal scaling can even be small enough to render direct application of a Krylov-subspace based iterative solver feasible. Because only cut elements of a specific shape cause small eigenvalues for systems based on maximum continuity splines in conjunction with diagonal scaling, the overall smallest volume fraction and

the scaled condition number are in general only weakly correlated in such systems [152].

3.2 Preconditioning and Literature Overview

Based on the conditioning analysis presented above, ill-conditioning of immersed FEM systems can effectively be negated by dedicated preconditioning techniques. The general idea of preconditioning is to construct a preconditioning matrix, \mathbf{B} , which is an approximation of the inverse of the system matrix \mathbf{A} , and then to solve the preconditioned system

$$\mathbf{B}\mathbf{A}\mathbf{x} = \mathbf{B}\mathbf{b} \quad (3.9)$$

Although the original system matrix \mathbf{A} in equation (2.8) can be ill-conditioned, i.e., $\kappa(\mathbf{A}) \gg 1$, a properly formed preconditioner results in a well-conditioned preconditioned system matrix, i.e., $\kappa(\mathbf{B}\mathbf{A}) \approx \kappa(\mathbf{A}^{-1}\mathbf{A}) = \kappa(\mathbf{I}) = 1$. In constructing the preconditioner \mathbf{B} , one balances the computational effort required to compute and apply the preconditioner with the extent to which the inverse of the original system matrix is approximated.

Various dedicated preconditioning techniques to resolve ill-conditioning problems caused by small cut elements have been developed in the context of GFEM and XFEM, such as: a preconditioner based on local Cholesky decompositions [164], a FETI-type preconditioner tailored to XFEM [165], an algebraic multigrid preconditioner that is based on the Schur complement of the enriched basis functions [166], and domain decomposition preconditioners based on additive Schwarz [167, 168].

In recent years, dedicated preconditioners have been developed for immersed FEM. It is demonstrated in [169] that systems with linear bases can effectively be treated by a diagonal preconditioner in combination with the removal of very small basis functions. Under certain restrictions on the cut-element geometry, it is derived in [170] that a scalable preconditioner for linear bases is obtained by combining a Jacobi preconditioner for basis functions on cut elements with a standard multigrid preconditioner for interior basis functions that do not intersect the boundary. In [152], a preconditioner is developed that combines diagonal scaling with local Cholesky factorizations. This technique is motivated by the analysis of the conditioning problems of immersed methods in the previous section, and can be interpreted as a local change of basis on small cut elements, where the Cholesky factorizations correspond to local orthonormalization procedures for almost linearly dependent basis functions, which are identified by a tailored algorithm. The resulting preconditioner effectively resolves ill-conditioning for immersed methods discretized with higher-order (continuous) bases.

Recently, Schwarz preconditioners—which were already discussed above in the context of GFEM and XFEM [167, 168]—have also gained momentum in immersed finite elements. The concept of Schwarz preconditioning to overcome the problem of linear dependencies on small cut elements was considered in [64, 171] for both linear and higher-order discretizations. This concept was generalized to multilevel hp -finite element bases in [63], where it was also demonstrated to be effective in parallel computing frameworks. In [172], Schwarz preconditioning of unfitted systems was applied as a smoother in a multigrid solver, making it suitable for large scale computations. The methodology was applied in high-performance parallel-computing settings in [173, 174], with [174] demonstrating excellent scalability for problems with multi-billion degrees of freedom distributed over close to 10^5 cores. A Balancing Domain Decomposition by Constraints (BDDC) scalable preconditioner, tailored to immersed FEM by choosing appropriate weighting coefficients for cut basis functions, has been proposed in [175]. BDDC methods are multilevel additive domain decomposition algorithms that can scale up to millions of cores/subdomains [176]. The preconditioner in [175] results in an effective preconditioner for linear basis functions and exhibits the same parallelism potential as standard BDDC, thus being well-suited for large-scale systems on distributed-memory machines.

It is worth mentioning that dedicated preconditioners have also been developed and investigated for ghost-penalty stabilized discretizations. While, as will be demonstrated in Sect. 4, such stabilized systems do not suffer from the typical conditioning problems related to small cut elements discussed in Sect. 3.1, the different setting of the problem and the involvement of additional terms does warrant the investigation into the applicability of efficient solvers, in particular multigrid techniques. Dedicated multigrid routines for Nitsche-based ghost-penalty stabilized methods are presented in [177–179]. It is notable that [179] additionally presents multigrid routines for Lagrange multiplier based methods, see also [180, 181]. Furthermore, in [182] and [183] a similar approach as in [170] (discussed above) is followed. In these references it is demonstrated that a stable splitting of the degrees of freedom exists, and that a scalable solver is obtained by combining a standard multigrid technique for a well-defined set of internal degrees of freedom with a diagonal preconditioner for the set of degrees of freedom along the boundary. While this technique is only applicable to ghost-penalty stabilized systems, in contrast to [170] it can also be applied to higher-order discretizations. A final noteworthy contribution is [184], which presents a preconditioner for ghost-penalty stabilized immersed interface problems of high contrast. In this reference a different splitting of

the degrees of freedom is applied to define a Schur complement, based on which preconditioners are presented that are robust to high contrast ratios.

The remainder of this section focuses on Schwarz preconditioning and the considerations regarding its application to systems derived from immersed finite element methods.

3.3 Schwarz Preconditioning

3.3.1 Concept of Schwarz Preconditioning

The concept of Schwarz preconditioning is to invert (restrictions of) local blocks of the system matrix $\mathbf{A} \in \mathbb{R}^{N \times N}$ and then sum these contributions to form the preconditioner. To provide a definition, we consider a set of (potentially overlapping) index blocks, where each index block contains $M_i \geq 1$ indices. The additive Schwarz preconditioner is then defined as

$$\mathbf{B}_{AS} = \sum_i \mathbf{B}_i = \sum_i \mathbf{P}_i \underbrace{(\mathbf{P}_i^T \mathbf{A} \mathbf{P}_i)^{-1}}_{\mathbf{A}_i^{-1}} \mathbf{P}_i^T \quad (3.10)$$

and the multiplicative Schwarz preconditioner as

$$\mathbf{B}_{MS} = \sum_i \mathbf{B}_i \prod_{j=1}^{i-1} (\mathbf{I} - \mathbf{A} \mathbf{B}_j) \quad (3.11)$$

The prolongation operator $\mathbf{P}_i \in \mathbb{R}^{N \times M_i}$ consists of the unit vectors corresponding to the M_i indices in the i -th index block. Pre- and post-multiplying the system matrix $\mathbf{A} \in \mathbb{R}^{N \times N}$ with the (transpose of) this prolongation operator restricts it to the submatrix $\mathbf{A}_i \in \mathbb{R}^{M_i \times M_i}$ consisting of only the indices in the i -th index block. Similarly, the opposite pre- and post-multiplication with these operators injects the local inverse $\mathbf{A}_i^{-1} \in \mathbb{R}^{M_i \times M_i}$ into the matrix $\mathbf{B}_i \in \mathbb{R}^{N \times N}$. It is to be noted that the index blocks may overlap. In additive Schwarz these contributions are then added in \mathbf{B}_{AS} and treated simultaneously. In multiplicative Schwarz, repetitions of the same index are treated sequentially. Application of the multiplicative Schwarz preconditioner can be expressed by recursive relations; see, e.g., [172]. It is noted that index blocks consisting of a single index simply reduce to Jacobi preconditioning in additive Schwarz, and to Gauss-Seidel preconditioning in multiplicative Schwarz. The formal definition and details about the construction of Schwarz preconditioners are presented [64] and [172].

3.3.2 Application to Immersed Systems

Schwarz preconditioners can be conceived of as locally orthonormalizing the basis functions corresponding to the indices in a block. In order to effectively employ the concept

of Schwarz preconditioning as a tailored preconditioner for unfitted systems, it is therefore essential that *for every set of almost linearly dependent functions, there is an index block containing all these functions*. Since almost linear dependencies occur between basis functions that are supported on a small cut element, the index blocks are generally chosen by selection procedures based on the overlapping support of basis functions, either for all active elements in \mathcal{T}_h or only for the cut elements in $\mathcal{T}_h^{\text{cut}}$. A discussion regarding considerations in the index blocks is provided later in this section.

A further interpretation of Schwarz preconditioning in relation to immersed finite element methods can be obtained from the additive Schwarz lemma. This lemma states that for a Symmetric Positive Definite (SPD) matrix \mathbf{A} , the $\mathbf{B}_{\text{AS}}^{-1}$ -inner product of an arbitrary vector \mathbf{y} is equal to [185, 186] (see [187, 188] for this specific form)

$$\begin{aligned} \mathbf{y}^T \mathbf{B}_{\text{AS}}^{-1} \mathbf{y} &= \min_{\mathbf{y} = \sum_j \mathbf{P}_j \tilde{\mathbf{y}}_j} \sum_i \tilde{\mathbf{y}}_i^T \mathbf{A}_i \tilde{\mathbf{y}}_i \\ &= \min_{\mathbf{y} = \sum_j \mathbf{P}_j \tilde{\mathbf{y}}_j} \sum_i (\mathbf{P}_i \tilde{\mathbf{y}}_i)^T \mathbf{A} (\mathbf{P}_i \tilde{\mathbf{y}}_i) \end{aligned} \quad (3.12)$$

In these identities, $\tilde{\mathbf{y}}_i \in \mathbb{R}^{M_i}$ denotes a block vector corresponding to the i -th index block. The statement $\mathbf{y} = \sum_j \mathbf{P}_j \tilde{\mathbf{y}}_j$ indicates that the sum of the prolongations of these block vectors form a partition of the vector $\mathbf{y} \in \mathbb{R}^N$. A set of block vectors with this property exists, provided that every index is contained in (at least) one index block. In the case that the blocks do not overlap, this set of block vectors is unique. In the case that the index blocks do overlap, multiple sets of block vectors have the partition property. Accordingly, the lemma states that the $\mathbf{B}_{\text{AS}}^{-1}$ -inner product of \mathbf{y} is equal to the minimum of the sum of the \mathbf{A}_i -inner products of the block vectors $\tilde{\mathbf{y}}_i$ over all sets of block vectors with the partition property.

To relate this to the specific conditioning problems in immersed finite element methods, recall that \mathbf{B}_{AS} can be considered as a sparse approximation of \mathbf{A}^{-1} . Efficient preconditioning requires \mathbf{B}_{AS} to have similar properties as \mathbf{A}^{-1} or, similarly, $\mathbf{B}_{\text{AS}}^{-1}$ to have similar properties as \mathbf{A} . The analysis of the conditioning problems associated with immersed finite element methods in Sect. 3.1 conveys that the principal cause of these problems is that, potentially, $\|\mathbf{y}\|_2 \gg \|\mathbf{y}\|_{\mathbf{A}}$ in case \mathbf{y} corresponds to a function comprised of (i) a very small basis function or (ii) almost linearly dependent basis functions. For the first case, it follows from (3.12) that \mathbf{y} must also have a small $\mathbf{B}_{\text{AS}}^{-1}$ -inner product, such that this property is captured by the additive Schwarz preconditioner, independent of the index blocks. For the second case, assume that indeed *for every set of almost linearly dependent functions, there is an index block containing all these functions*, in accordance with the previously stated postulate. Then, in case \mathbf{y} corresponds to a function comprised of

almost linearly dependent basis functions, \mathbf{y} can be written as the prolongation of a single block vector. Consequently, also in the second case, it follows from (3.12) that \mathbf{y} will have a small $\mathbf{B}_{\text{AS}}^{-1}$ -inner product, such that this property is captured in the additive Schwarz preconditioner. As a result, with a proper choice of the index blocks, for both causes of very small eigenvalues of the matrix \mathbf{A} , the corresponding modes will also be present in $\mathbf{B}_{\text{AS}}^{-1}$. The Schwarz preconditioner thus specifically targets the problematic aspects of small eigenvalues due to small cut elements (for both linear and higher-order discretizations), and thereby effectively resolves the ill-conditioning in systems derived from immersed formulations.

3.3.3 Effectivity and Multigrid Preconditioning

While the effectivity of Schwarz preconditioning for immersed problems can be explained by the additive Schwarz lemma in combination with a particular selection of the blocks, a formal mathematical bound on the eigenvalues of an unfitted system treated by Schwarz preconditioning has not yet been formulated. Numerical results, however, consistently show that the resulting systems behave the same as boundary-fitted systems, in the sense that (for second order problems as the Poisson equation) the condition number scales as h^{-2} , and the number of iterations that is required to solve the linear system up to a prescribed tolerance is proportional to h^{-1} [64, 172].

To resolve the remaining grid dependence after Schwarz preconditioning, the Schwarz preconditioner can be applied as a smoother in a geometric multigrid framework. In [172–174] this is demonstrated to result in a methodology that is robust to cut elements and which solves linear systems with quasi-optimal complexity, i.e., at a computational cost that is linear with the number of degrees of freedom (DOFs). A delicate consideration in a multigrid framework is the choice between additive and multiplicative Schwarz. As demonstrated in [172], the stability of a multigrid solver with an additive Schwarz smoother requires a considerable amount of relaxation, and for this reason [172] employs a multiplicative implementation. As discussed later in this section, multiplicative Schwarz is less suited for parallelization, such that parallel implementations employ either additive Schwarz [174] or a hybrid variant [173]. Another important aspect of multigrid solvers with Schwarz smoothers is the dependence of their effectivity on the discretization order, which is investigated and discussed in [172].

3.3.4 Block Selection

As previously mentioned, for Schwarz preconditioning to be effective, it is essential that every combination of basis functions that can become almost linearly dependent is

contained in an index block, which is generally achieved by selecting these blocks based on the overlap in the supports of basis functions. In [64], an index block is devised for every cut element, containing all basis functions supported on it. For basis functions that are not supported on any cut element, a simple diagonal scaling is performed. A tailored block selection procedure for multilevel hp -adaptive discretizations is developed in [63]. In this procedure, the set of basis functions supported on a refined (leaf) element is further restricted to only the necessary DOFs. Additionally, the procedure is optimized by only devising blocks for cut elements with a volume fraction that is smaller than a certain threshold $\eta^* \in [0, 1]$. For locally refined discretizations based on truncated hierarchical B-splines, a block selection strategy is presented in [172]. In [64] it is noted that, for vector-valued problems, degrees of freedom describing the solution in a certain geometrical dimension can generally not form a linear dependency with degrees of freedom describing the solution in another geometrical dimension. Hence, it can be beneficial to generate separate index blocks for each geometrical dimension.

The most important consideration in the selection of index blocks is the size of the blocks. Small blocks are computationally inexpensive, but can miss almost linear dependencies on pathologically cut elements. Large blocks are more robust, but are computationally more expensive. With the Schwarz preconditioner directly employed in an iterative solver, the number of iterations is generally large enough for small blocks to be more efficient. When the Schwarz procedure is employed as a smoother in a multigrid method, the small number of iterations generally renders larger blocks more appropriate. For this reason, [172] and [174] consider block selection procedures based on multiple elements. In [172] an index block is created for every basis function, containing all the basis functions with a support that is encapsulated by the support of the function for which the block is created. Ref. [174] employs index blocks containing all basis functions supported on a cluster of 2^d elements, with d the number of dimensions.

Regarding the lack of consensus on the block selection procedures in different contributions, it can be concluded that an optimal choice of index blocks is still an unresolved question.

3.3.5 Implementation Aspects

A specific operation in the construction of a Schwarz preconditioner is the computation of stable inverses of the submatrices \mathbf{A}_i . With very small eigenvalues, numerical round-off errors can cause detrimental errors in the inversion of submatrices with eigenvalues that are too close to machine precision. In the worst case, this can lead to negative

eigenvalues in the preconditioned system, resulting in failure of the iterative solver. For that reason, the inverses are generally not computed by a simple inversion operation, but via an eigenvalue decomposition. After the decomposition, eigenvalues that are smaller than a prescribed threshold are discarded. Details about this procedure are described in [63, 64].

The Schwarz preconditioner is suitable for parallel implementations [63]. A particular facet of the parallelization of a Schwarz preconditioner is that each submatrix \mathbf{A}_i has to be available in a single subprocess for inversion. As this is not generally the case in parallel (boundary-fitted) finite element codes, this calls for special care in the implementation, e.g., by applying ghost elements [63]. As previously mentioned, a specific consideration pertaining to parallel multigrid implementations is the choice between additive and multiplicative Schwarz. While multiplicative Schwarz does not require relaxation, it does require extensive communication and synchronization between parallel processes, specifically in distributed memory systems. For that reason the parallel multigrid implementation in [174] employs an additive Schwarz smoother, while in [173] a hybrid variant is applied with an additive approach for DOFs that are shared between processes and a multiplicative treatment of DOFs that belong to a single process.

3.3.6 Extension to Other Problems

The additive Schwarz lemma in (3.12) only pertains to Symmetric Positive Definite (SPD) systems. Such systems cover a large number of applications, encompassing many problems in structural mechanics, but do not comprise all problems commonly solved by immersed finite element methods. In particular, incompressible-flow problems in which the pressure takes the form of a Lagrange multiplier in the weak formulation are indefinite, and problems involving convection or formulations based on the nonsymmetric Nitsche method are not symmetric. Applications of the Schwarz framework to immersed finite element approximations of such problems are considered in [64], in which for all considered problems it is observed that the solver is independent of the cut elements and that the number of iterations is approximately inversely proportional to the mesh size.

4 Stability

4.1 Stability Analysis

As indicated in Sect. 2.2.3, small cut elements do not only adversely affect the conditioning of the resulting linear system, but can also disturb the quality of the solution as

the normal gradient is not adequately controlled. When Nitsche's method is applied on a boundary-fitted discretization, the trace inverse inequality $\|\partial_n u_h\|_{T \cap \partial\Omega_D} \lesssim h_T^{-1/2} \|\nabla u_h\|_T$ holds on every element T adjacent to the Dirichlet boundary $\partial\Omega_D$; see, e.g., [100]. By virtue of this trace inequality, the bilinear operator $a_h(\cdot, \cdot)$ defined in (2.7) is coercive with a global Nitsche parameter β that scales with h^{-1} , and optimal error bounds with respect to the $H^1(\Omega)$ -norm hold. With an immersed discretization, however, an element T that intersects $\partial\Omega_D$ only partially intersects the physical domain Ω . Therefore, the aforementioned trace inverse inequality assumes the form $\|\partial_n u_h\|_{T \cap \partial\Omega_D} \lesssim h_{T_\Omega}^{-1/2} \|\nabla u_h\|_{T_\Omega}$, where h_{T_Ω} indicates a generalized thickness of the fragment $T_\Omega = T \cap \Omega$ (the intersection of element T and the physical domain) normal to $T \cap \partial\Omega_D$ (the intersection of element T and the Dirichlet boundary); see Sect. 2.1. Consequently, coercivity of the bilinear form $a_h(\cdot, \cdot)$ requires an element-wise Nitsche parameter β that is not inversely proportional to the characteristic mesh width of the background element, h_T , but instead is inversely proportional to the generalized thickness of the intersection, i.e., $\beta \sim h_{T_\Omega}^{-1} \gtrsim h_T^{-1}$. Due to the fact that without special precautions, in immersed finite element methods one has no control over h_{T_Ω} , the Nitsche parameter can in principle become arbitrarily large.

In [1] it is shown that the error of immersed Nitsche-based formulations admits a natural analysis in the norm (within this manuscript referred to as the β -norm)

$$\|u\|_\beta^2 = \|\nabla u\|_\Omega^2 + \|\beta^{1/2} u\|_{\partial\Omega_D}^2 + \|\beta^{-1/2} \partial_n u\|_{\partial\Omega_D}^2 \quad (4.1)$$

$$u \in V_h \oplus H^2(\Omega)$$

provided that β is large enough for the bilinear form to be coercive. Note that this norm is referred to differently in [1], and that it differs from the energy norm that will be introduced in Sect. 4.5. On account of the last term, this norm cannot be applied to the full space $H^1(\Omega)$. For this reason, it is assumed in the analysis that $u \in H^2(\Omega)$, and the setting is restricted to the composite space $V_h \oplus H^2(\Omega)$. In [1] it is demonstrated that the Nitsche-based approximation possesses a best-approximation property with respect to the β -norm, but that the error of the best approximation in this norm can still be arbitrarily large. In view of this and the fact that the equivalence between the β -norm in (4.1) and the $H^1(\Omega)$ -norm can be arbitrarily weak, error bounds in the $H^1(\Omega)$ -norm cannot be provided. Ref. [1] provides examples of computations where direct application of Nitsche's method fails for unfitted problems and provides references in which nonphysical stress patterns on small cut elements are reported.

4.2 Remedies and Literature Overview

Multiple stability-enhancing techniques have been developed to overcome the aforementioned problem, all of which are applicable to higher-order discretizations. A list of these is presented below. Two techniques in particular are discussed in detail in the subsequent subsections, *viz* element aggregation in Sect. 4.3 and ghost-penalty stabilization in Sect. 4.4. Essentially, these approaches provide control over the gradients on cut elements in terms of the gradients on interior elements, in a manner that is consistent with the original problem. Consequently, extended coercivity holds with respect to $H^1(\Omega_h)$ (*n.b* on the entire active domain), and the bilinear form is coercive with a well-behaved Nitsche parameter $\beta \sim h_T^{-1} (\lesssim h_{T_\Omega}^{-1})$. With element aggregation this is achieved through a modification of the approximation space, while with the ghost penalty a consistent stiffness is added to the bilinear form. As a result, these techniques provide optimal approximation properties, and condition number estimates analogous to those for boundary-fitted finite element formulations. Therefore, both aspects of the small-cut-element problem are resolved simultaneously. A unified analysis of the two methods is presented in Sect. 4.5.

- By employing a *different formulation* to enforce boundary conditions, the stability problems caused by Nitsche's method are circumvented. Such alternative methodologies to enforce boundary conditions are discussed in Sect. 2.2.2 (e.g., the nonsymmetric Nitsche method [142–144] or the shifted boundary method [89–92]). For some applications, however, these approaches lead to other complications, such as loss of consistency, symmetry, or adjoint consistency.
- A *fictitious domain stiffness* is commonly employed with high-order discretizations in the finite cell method [2, 74] (see [75, 106] for reviews). This approach is mathematically analyzed in [189]. With a fictitious domain stiffness, the volume integrals of the bilinear operator are extended into the fictitious domain multiplied by a small parameter (in practical applications this only holds for the fictitious parts of cut elements, i.e., $\Omega_h \setminus \Omega$). This provides a stiffness on the fictitious part of cut elements, providing some control over the approximate solution on these. This approach has been applied to many real-world applications, notable examples of which are implant-vertebra models [21] and additively manufactured structures [103]. The fictitious domain stiffness is not consistent with the original problem, however. For this reason, a drawback of this approach is that it can be challenging (or

for certain applications impossible) to set the parameter low enough to keep the consistency error at an acceptable level, but high enough to obtain sufficient control over the fictitious parts of elements. Besides the stability, the fictitious domain stiffness also improves the conditioning, but generally not to a level that permits the application of iterative solvers.

- A *minimal stabilization procedure* for Nitsche's method is presented in [190], which is inspired by the technique developed for a Lagrange multiplier method in [191]. In this procedure, the elements are first classified as good (resp. bad) elements, i.e., elements with a sufficiently (resp. insufficiently) large overlap with the physical domain. Next, for each bad element, the normal derivative of the discrete function on the Dirichlet boundary that appears in the boundary integral of the bilinear form, is replaced by the normal derivative of an extension of the function on a nearby good element. This gives rise to a similar trace theorem as in boundary-fitted discretizations and, accordingly, the weak formulation is coercive with a Nitsche parameter that is inversely proportional to the (untrimmed) element size of the background grid such that optimal error estimates can be derived. This technique has mainly been applied to trimmed patches in isogeometric analysis, and can also be employed for coupling conditions on unfitted interfaces in overlapping multi-patch discretizations [192]. Additionally, in [193] it is shown that this methodology can be used to obtain inf-sup stable mixed formulations. It is to be noted that this approach does not provide control of the gradients in the (fictitious part of) cut elements themselves, and therefore only provides coercivity with respect to $H^1(\Omega)$, not extended coercivity with respect to $H^1(\Omega_h)$. As such, this approach yields stability and provides a well-posed imposition of boundary conditions, but it does not solve conditioning problems.
- *Recovery-based stabilization techniques* [194, 195] are commonly applied in the Cartesian grid finite element method (cgFEM, [57, 82]). In [194] a technique similar to the minimal stabilization procedure discussed above is presented, and [195] weakly constrains small cut elements with an L^2 -projection terms (over the entire background elements), similar to the ghost penalty with patch control discussed later in this section. These iterative techniques do not employ extensions of the (current) u_h and v_h on nearby elements, however, but instead consider smooth extensions from a previous iteration (hence, recovery-based), and Richardson iteration is applied to obtain the final result. As a benefit of this technique, [195] mentions that it does not require operators that are not generally available in (immersed) finite element codes. Furthermore, [194] and [195] establish the (mesh) convergence of the solution and the convergence of the

Richardson iteration and [195] also establishes a bound on the condition number similar to that of boundary-fitted finite element methods.

- A *least squares stabilization* term can be applied when the employed function space is (at least) C^1 -continuous. In [196] and [197] it is demonstrated that a stable and coercive formulation can also be achieved by applying a least squares finite element term in the vicinity of the Dirichlet boundary. Because in $H^2(\Omega)$ the normal gradient on the boundary can be controlled by volumetric terms, coercivity is even achieved in the full space instead of only in the discrete space. A drawback of this approach is that it only applies to C^1 -continuous bases. The least squares stabilization does not repair the conditioning, such that in [196] this approach is combined with the removal of basis functions with very small supports in the physical domain and in [197] it is combined with a fictitious domain stiffness.

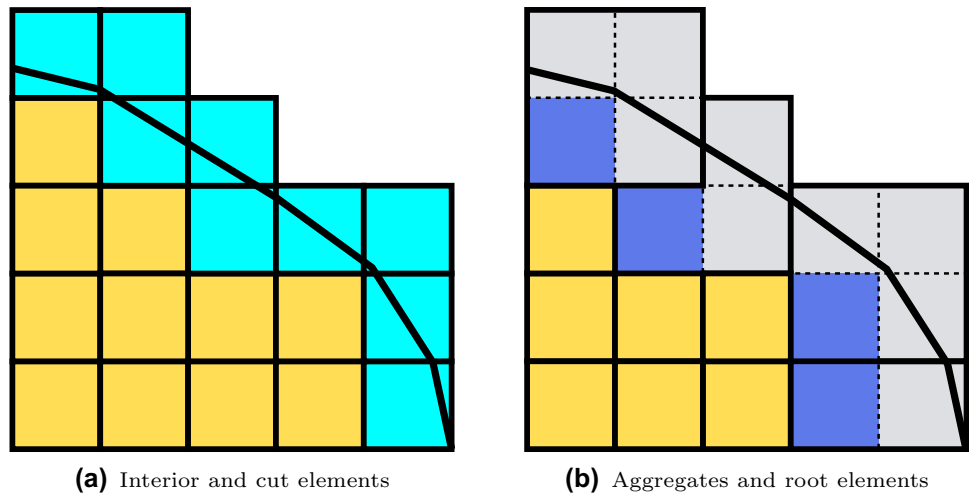
4.3 Aggregated Finite Element Methods

In this section, we introduce a methodology to solve both the stability and conditioning issues described in Sect. 2.2.3. The method is based on an aggregation framework proposed in [80] and is coined the *aggregated* finite element method (AgFEM). As the name points out, the method relies on an aggregation (also called agglomeration) of elements. This element aggregation is used to define a discrete extension operator from degrees of freedom on well-posed (interior) elements to ill-posed (cut) elements. The AgFEM space is the image of this operator. Thus, as indicated in Sect. 4.2, AgFEM solves the issues described above by modifying the finite element space, without altering or perturbing the bilinear form. We present this idea for (nodal or Lagrangian) C^0 -continuous finite element spaces. It can also be applied to discontinuous Galerkin methods, even though this case is trivial; discontinuous Galerkin schemes can readily be applied to the polytopal meshes obtained after aggregation. Similar methods have been described in [37, 83–85, 88, 153, 198–202].

4.3.1 Geometrical Construction

Let us recall from Sect. 2.1 the problem domain $\Omega \subset \mathbb{R}^d$ that is embedded in the ambient domain $\mathcal{A} \supset \Omega$. The mesh $\mathcal{T}_h^{\mathcal{A}}$ is a quasi-uniform and shape-regular discretization of \mathcal{A} , \mathcal{T}_h is the set of all active elements $T \in \mathcal{T}_h^{\mathcal{A}}$ that intersect Ω , and the union of all active elements is defined as $\Omega_h = \cup_{T \in \mathcal{T}_h} T \supset \Omega$. The set of cut elements that intersect the boundary $\partial\Omega$ of the problem domain is denoted as $\mathcal{T}_h^{\text{cut}} \subset \mathcal{T}_h$ and additionally we define the set of interior elements as $\mathcal{T}_h^{\text{in}} = \mathcal{T}_h \setminus \mathcal{T}_h^{\text{cut}}$, see Fig. 6a. As previously mentioned, immersed finite element formulations can lead to ill-conditioned discrete systems and

Fig. 6 On the left, we show the active portion of a background mesh \mathcal{T}_h , with interior elements in $\mathcal{T}_h^{\text{in}}$ (yellow) and cut elements in $\mathcal{T}_h^{\text{cut}}$ (light blue). On the right, we create aggregates (in gray) composed of one root element (dark blue) and some cut elements. The curved black line represents the boundary $\partial\Omega$



unbounded gradients on cut elements. The interior elements in $\mathcal{T}_h^{\text{in}}$ do not play a role in these problems, but it is the cut elements in $\mathcal{T}_h^{\text{cut}}$ that can lead to the so-called small-cut-element problem.

The definition of finite element spaces that are robust to cut-element configurations makes use of an element aggregation strategy. The idea is to aggregate cut elements to interior elements. An aggregate contains one interior element (the root element) and one or more cut elements. We do not make any assumption on the shape of aggregates. As a result of this finite element space definition, aggregation is only active on the boundary. Interior elements that are not in touch with cut elements are not affected. We represent the aggregated mesh with $\mathcal{T}_h^{\text{ag}}$. In Fig. 6b, cut elements have been aggregated to interior elements, creating aggregates. The root element is highlighted with a different color in each aggregate. The aggregation algorithm is straightforward and can be found, e.g., in [80]. The algorithm proposed in this reference is iterative. At each iteration, cut elements are aggregated to an interior element or to a previously aggregated element (after the first iteration). The aggregation should minimize the aggregate size, since this will affect the final accuracy of the solution.

Remark 4.1 In this discussion, we considered for simplicity that cut elements are ill-posed and must be aggregated to interior elements. However, this naive definition of ill-posed elements is not a requirement of the method and more subtle definitions can be considered. For instance, one can define a threshold volume fraction $\eta^* \in [0, 1]$ (see (3.4)) and mark any element $T_i \in \mathcal{T}_h^{\text{cut}}$ with $\eta_i \leq \eta^*$ as ill-posed. The naive case is recovered for $\eta^* = 1$ and the standard Galerkin method with $\eta^* = 0$. One can verify that the condition number and stability bounds hold for this relaxed definition with constants that depend on η^* . $\eta^* = 1$ is an excellent choice

as soon as one wants to have enough resolution to capture geometrical details, which is the situation in most cases. For under-resolved situations, e.g., for thin-walled structures in which the thickness scale t_ϵ is not captured by the background mesh, a more clever choice of η^* , e.g., $\eta^* \sim \frac{t_\epsilon}{h}$, is required. With a proper choice of η^* , only elements small enough to affect the stability are aggregated, such that the effective mesh size at the boundary is not affected. Furthermore, the aggregate sizes can be locally reduced using adaptation to accurately capture complex geometrical features [203], which is discussed below in the Subsection Adaptive meshes. In practice, the choice of η^* is a trade-off between accuracy (aggregate size) and well-posedness (condition number bound) of the resulting linear system. We note that the definition and implementation of AgFEM is independent of η^* , since this parameter only affects the geometrical algorithm that aggregates elements.

4.3.2 Discontinuous Spaces

Let us consider first the case of discontinuous Galerkin finite element spaces. We denote by V_h^- a discontinuous Galerkin space of a given order. The main reason why standard finite element methods on unfitted meshes fail is that one has no control over $\|u_h\|_{\Omega_h}$. Control over this quantity implies the stability of the method regardless of the cut location. Let us consider a discontinuous Galerkin space on top of Ω_h (see Fig. 7a). The support of a shape function on a cut element $T \in \mathcal{T}_h^{\text{cut}}$ is $T \cap \Omega$. This intersection depends on the cut location, and its measure can be arbitrarily small. In the limit of vanishing measure, the value of the degree of freedom associated with this shape function does not affect the linear system. Thus, the problem is singular.

Given a mesh \mathcal{T}_h , we can define the discontinuous element-wise spaces

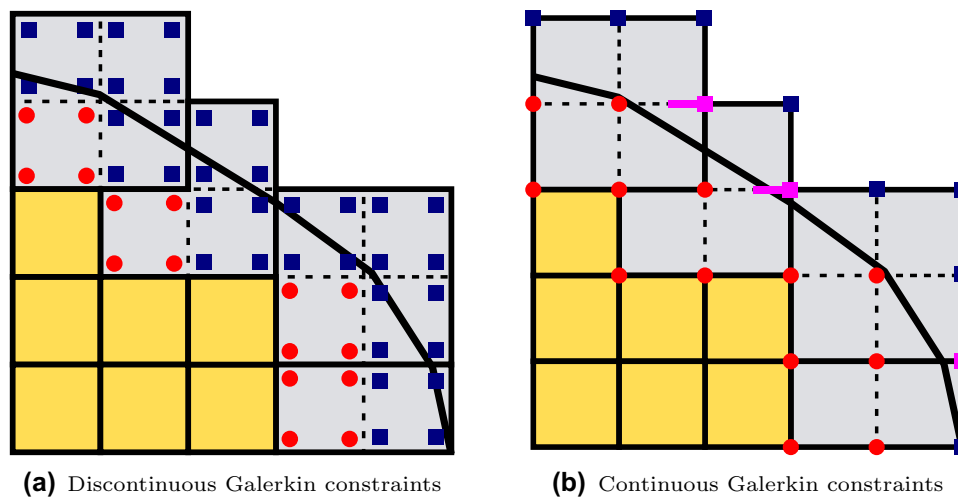


Fig. 7 On the left, we depict the degrees of freedom of a piecewise linear discontinuous Galerkin space $\mathcal{Q}_1^-(\Omega_h)$. The nodes are deliberately placed in the element interiors to make it evident that they do belong to the element. Red circles represent well-posed degrees of freedom and blue squares ill-posed ones. In this case, for each ill-posed degree of freedom, its constraining degrees of freedom are the well-posed ones in the same aggregate using the expression in (4.4). On the right, we glue together degrees of freedom between elements

to enforce C^0 continuity. In order to define the constraints, we must provide the concept of the aggregate owner of ill-posed degrees of freedom. This is obvious for aggregate-interior nodes (there is only one option). For the ill-posed degrees of freedom on aggregate interfaces, lines point to the owner aggregate. With this ownership information, one can now constrain the ill-posed degrees of freedom by the well-posed ones in the aggregate that owns them using the same expression (4.4) as above

$$\mathcal{P}_p^-(\mathcal{T}_h) = \prod_{T \in \mathcal{T}_h} \mathcal{P}_p(T), \quad \mathcal{Q}_p^-(\mathcal{T}_h) = \prod_{T \in \mathcal{T}_h} \mathcal{Q}_p(T)$$

depending on the local finite element space, namely polynomials up to order p or the tensor product of univariate polynomials of order p . The discontinuous Galerkin space V_h^- on the active mesh can either be $\mathcal{P}_p^-(\Omega_h)$ or $\mathcal{Q}_p^-(\Omega_h)$. Other spaces, like serendipity finite elements, could readily be considered. The space in Fig. 7a corresponds to $\mathcal{Q}_1^-(\Omega_h)$.

Based on the previous discussion, it is straightforward to check that degrees of freedom on interior elements are not problematic. The corresponding shape function has support on a whole element. Only the degrees of freedom on cut elements are problematic. Discontinuous Galerkin methods are geometrically flexible, i.e., the functional space and the topology of the elements are not connected. Thus, we can simply define the aggregated discontinuous Galerkin space as $\mathcal{P}_p^-(\mathcal{T}_h^{\text{ag}})$ or $\mathcal{Q}_p^-(\mathcal{T}_h^{\text{ag}})$. All these shape functions are polynomials, and their support is at least one full interior element. Thus, it can be checked that

$$\|v_h\|_{\Omega_h} \lesssim \|v_h\|_{\Omega} \quad v_h \in V_h^{\text{ag},-} \quad (4.2)$$

Extended stability of gradients is obtained using the same argument. Gradients of functions in $V_h^{\text{ag},-}$ are polynomials of degree $p-1$ in each aggregate, and the norm on the whole aggregate can be bounded by the one on the root element, getting

$$\|\nabla v_h\|_{\Omega_h} \lesssim \|\nabla v_h\|_{\Omega} \quad v_h \in V_h^{\text{ag},-} \quad (4.3)$$

These relations are also employed in the mathematical analysis of stabilized methods in Sect. 4.5. The constants in the extended stability bounds (4.2) and (4.3) depend on the order of approximation, but are independent of the background mesh size h and the cut location. Thus, the small-cut-element problem is solved when using discontinuous Galerkin methods on aggregated meshes. Similarly, cell merging strategies in the finite volume setting have been proposed, e.g., in [204]. Its application to the discontinuous Galerkin method can be found in [86, 122].

The aggregation method preserves the order of accuracy if the aggregate size is proportional to the background mesh size. The ratio between the aggregate size and root element size is related to the geometry of the domain boundary (see [80] for more details) and can be improved with refinement (for a given definition of the domain boundary $\partial\Omega$). For a fixed mesh representation of the domain boundary (the standard situation), the assumption that the aggregate size is proportional to the background mesh size holds for h small enough.

4.3.3 Constraining Discontinuous Spaces

It is straightforward to see that $V_h^{\text{ag},-} \subset V_h^-$. Even though it is not really needed for the discontinuous Galerkin method, one can build $V_h^{\text{ag},-}$ by constraining V_h^- . The idea is to define

an extension operator $\mathcal{E}_h^{\text{ag},-} : V_h^{\text{in},-} \rightarrow V_h^{\text{ag},-}$, with $V_h^{\text{in},-}$ the restriction of V_h^- to $\mathcal{T}_h^{\text{in}}$. In the discontinuous Galerkin case, this extension operator can be defined aggregate-wise. The idea is to extend the root-element shape functions to the whole aggregate. This is equivalent to constraining degrees of freedom on cut elements to the degrees of freedom on the respective root elements.

Let us recall Ciarlet's definition of finite elements; see, e.g., [96]. For nodal finite element methods (continuous or discontinuous), one can define a set of points (so-called nodes) and an associated set of degrees of freedom, corresponding to pointwise evaluation at the nodes. Nodes are chosen such that the degrees of freedom uniquely define functions in V_h^- and thus define a basis for its dual space. The dual basis of the degrees of freedom are the so-called shape functions. Thus, there is a correspondence between nodes, degrees of freedom and shape functions. We use the following notation: given a node α_i , we represent its corresponding shape function and degree of freedom by ϕ_i and $\sigma_i(\cdot)$, respectively. The duality between the degrees of freedom and the shape functions implies $\sigma_i(\phi_j) = \delta_{ij}$ with δ_{ij} the Kronecker delta, and the relation between the nodes and the degrees of freedom in turn implies $\phi_i(\alpha_j) = \delta_{ij}$. By virtue of the dual relation between degrees of freedom and shape functions, the following correspondence holds for all $v_h^- \in V_h^-$: $y_i = \sigma_i(v_h^-)$ if and only if $v_h^- = \sum_i y_i \phi_i$.

We can classify the degrees of freedom as *well-posed* (the ones of the root element) and *ill-posed* (the ones on the cut elements). We denote the set of well-posed (resp., ill-posed) degrees of freedom by $\mathcal{O}_h^{\text{wpd}}$ (resp., $\mathcal{O}_h^{\text{ipd}}$). Next, we define a map $\mathcal{O}_h^{\text{ipd} \rightarrow \text{wpd}}(\cdot)$ that for an ill-posed degree of freedom returns the well-posed ones. For discontinuous Galerkin methods, this map is defined aggregate-wise. At each aggregate $T \in \mathcal{T}_h^{\text{ag}}$, the ill-posed degrees of freedom (blue squares in each aggregate of Fig. 7a) are constrained by the values of its corresponding well-posed degrees of freedom (the ones in the corresponding root element, red circles in Fig. 7a). More specifically, an ill-posed degree of freedom $\sigma_i(\cdot)$ is constrained as

$$\begin{aligned} \sigma_i(\cdot) &= \sum_{\sigma_j \in \mathcal{O}_h^{\text{ipd} \rightarrow \text{wpd}}(\sigma_i)} \sigma_j(\phi_j) \sigma_j(\cdot) \\ &= \sum_{\sigma_j \in \mathcal{O}_h^{\text{ipd} \rightarrow \text{wpd}}(\sigma_i)} \phi_j(\alpha_i) \sigma_j(\cdot) \end{aligned} \quad (4.4)$$

where the basis functions ϕ_j of well-posed degrees of freedom are evaluated outside of the elements on which these are defined in $\sigma_i(\phi_j)$ and $\phi_j(\alpha_i)$. Equation (4.4) introduces constraints on ill-posed values, stating that the values on ill-posed nodes are not free but are determined from the well-posed ones. More specifically, the ill-posed value is equal to the nodal values on the root element multiplied by

the value of the corresponding shape function evaluated at the ill-posed node. Doing this, we are constraining the ill-posed degrees of freedom to be the extension of the values in the interior. This extension applied to functions in V_h^- defines $\mathcal{E}_h^{\text{ag},-}$, whose image is $V_h^{\text{ag},-}$.

4.3.4 Aggregation for Continuous Spaces

Now, we would like to generalize the extension-based definition of the aggregated discontinuous finite element space to C^0 Lagrangian spaces. First, we discuss how the active and interior C^0 spaces can be obtained from the discontinuous space by enforcing C^0 continuity of piecewise polynomials using a local-to-global map. We also discuss why this approach cannot be easily applied to the aggregated meshes. Instead, we introduce the concept of ownership of ill-posed degrees of freedom on aggregates, and define an extension operator from well-posed to ill-posed degrees of freedom that preserves C^0 continuity.

We can readily define the C^0 Lagrangian finite element space as a subspace of the discontinuous space as $V_h = V_h^- \cap C^0(\Omega_h)$, and its restriction to interior elements is represented with V_h^{in} . In standard finite element methods, one can enforce C^0 continuity by using the previously introduced one-to-one relation between nodes, degrees of freedom, and shape functions. A shape function is different from zero only on its corresponding node, and zero on all the other nodes, such that the degree of freedom is just the evaluation on its corresponding node. C^0 continuity is enforced by the following equivalence class: two degrees of freedom must have the same value (irrespective of the element) if their corresponding nodes are in the same spatial point. Doing this, we glue together degrees of freedom from different elements. Figure 7b illustrates the result after gluing together the element-wise degrees of freedom in Fig. 7a. This enforcement of continuity is amenable for implementation: one defines an element-wise matrix assembly and a local-to-global index map for degrees of freedom.

Unlike discontinuous Galerkin methods, this construction only holds for particular element-wise polynomial spaces and element topologies. First, it is not straightforward to implement such continuity in aggregated meshes. Second, it can destroy the approximation properties of V_h^{ag} . Let us take a look at an ill-posed degree of freedom on the interface between two aggregates (purple nodes in Fig. 7b). In a straightforward extension of the constraints in (4.4), such a degree of freedom would be constrained by the multiple aggregates that contain it. This, together with the enforcement of C^0 -continuity, would couple the degrees of freedom between separate aggregates, and thereby could ruin approximation properties. In extreme cases, it can lead to

over-constrained systems. This *locking* phenomenon has been analyzed in detail in [81].

The aggregated finite element method in [80] was designed to solve this locking phenomenon and provide an accurate C^0 Lagrangian finite element space on unfitted meshes that is easy to implement. The idea is to introduce the concept of *ownership* for ill-posed degrees of freedom. For each ill-posed global degree of freedom in V_h^{ag} , one assigns one of the aggregates in $\mathcal{T}_h^{\text{ag}}$ containing this degree of freedom as the owner. The choice of the owner is arbitrary and does not affect the main properties of the method. We can define a map $\mathcal{O}_h^{\text{ipd} \rightarrow \text{wpd}}$ using the fact that each aggregate contains only one interior element. Given an ill-posed degree of freedom, its corresponding well-posed constraining degrees of freedom are those of the root element of the aggregate owner. We illustrate this construction in Fig. 7b. The purple elements are the ones that belong to more than one aggregate (the other ones are straightforward). For these degrees of freedom, the small rectangle points to the aggregate owner of the degree of freedom. With this construction, we can now constrain global ill-posed degrees of freedom using the same constraints as for discontinuous Galerkin methods (4.4).

These constraints are local and define an extension $\mathcal{E}_h^{\text{ag}} : V_h^{\text{in}} \rightarrow V_h^{\text{ag}}$. In fact, the aggregated finite element space is the image of this extension operator. The method satisfies the desired approximation properties and the extended $L^2(\Omega_h)$ and $H^1(\Omega_h)$ stability in (4.2) and (4.3); see [80] for more details. We also refer to [205] for an alternative definition of the discrete extension operator that relies on interpolation for high-order polynomial bases and to [200, 201, 206] for the extension of higher-order continuous (isogeometric or spline-based) discretizations. The definition of a space-time discrete extension operator for moving domains and interfaces can be found in [207].

4.3.5 Effectivity and Multigrid

One of the main motivations of aggregated finite element methods is to produce a well-posed discrete system. For second-order elliptic problems, one can prove the same condition number bounds as for boundary-fitted formulations. More specifically, the condition number bound is of the order of h^{-2} ; see [80, Corollary 5.9]. Thus, the condition number is independent of the mesh location. As a result, one can apply standard direct or iterative solvers and preconditioning techniques. However, the sparsity pattern of the matrix is different from standard finite elements. As in adaptive finite elements, the sparsity pattern is affected by the constraints.

The linear solver step in aggregated finite element methods has been studied in [208]. The authors propose a

row-wise linear algebra distribution layout for the resulting matrix in [208, §3.6] and minimize the amount of inter-processor communications. The authors use the parallel implementation of aggregated finite elements in FEMPAR [209]. They apply a standard algebraic multigrid preconditioner (called GAMG) from the PETSc library [210] using a standard configuration. The results in [208, §4] show excellent scalability and optimality properties for this standard solver applied to the aggregated finite element method. The aggregated results are very similar to the results obtained with boundary-fitted meshes. The largest problems that are considered exceed 10^9 elements on 16,464 processors. From these numerical experiments, it is observed that the additional coupling between boundary degrees of freedom does not affect a standard algebraic multigrid preconditioner.

4.3.6 Adaptive Meshes

The size of the aggregates on the boundary can harm the accuracy of the method. One way to limit the size of aggregates is via mesh refinement. However, local mesh refinement in continuous finite elements is not straightforward. In the resulting mesh so-called hanging nodes appear, which must be constrained by regular nodes to keep C^0 continuity. The nature of this constraint is analogous to the ones defined in (4.4) for aggregated finite elements.

The combination of mesh adaptivity and aggregation constraints is not straightforward. First, the definition of ill-posed versus well-posed degrees of freedom is more complex. A degree of freedom that only belongs to cut elements can be well-posed because it constrains well-posed hanging degrees of freedom. The right intertwining of these two sets of constraints has been analyzed in [203]. This work proposed a two-step algorithm to construct the discrete extension operator that carefully mixes aggregation constraints of problematic degrees of freedom and standard hanging degree of freedom constraints. Following this approach, the aggregated finite element space is defined as a discrete extension with well-defined linear constraints.

4.3.7 Implementation Aspects

The aggregation algorithm can readily be implemented in a finite element code that supports mesh adaptivity. First, local element matrices are computed as usual. Next, the constraints for the aggregated finite element method are direct constraints, i.e., constraints which are enforced in the local-to-global assembly by expressing the ill-posed degrees of freedom in terms of well-posed ones to eliminate them from the system.

The distributed-memory implementation of the aggregated finite element method has been considered in [208]. The proposed implementation starts with a standard

distributed implementation of the finite element method that relies on a sub-mesh partition (with gluing info, e.g., via one layer of neighbor elements). First, one must parallelize the aggregation algorithm. The parallel implementation of this step is straightforward; one needs to perform nearest neighbor communications at the end of each iteration of the aggregation algorithm to propagate information. The result is identical to the serial version (see [208, sect. 3.2]).

In general, the aggregation strategy is such that aggregates can have support on multiple processors. Thus, constraints in (4.4) cannot be computed locally. To compute these constraints in a parallel environment, one requires information from the root element. The cut element to root element map (and its processor Id) is a by-product of the parallel aggregation step. This way, it is easy to define the information that each processor must receive from other processors to compute constraints. However, the inverse map that provides the cut elements constrained by a given root element is not straightforward. This map is required to prepare the data that needs to be sent to other processors. The parallel inverse path reconstruction algorithm in [208, §3.5] generates this information.

4.3.8 Extension to Other Problems

The aggregated spaces can be used to discretize elliptic and parabolic problems (combined with some time-stepping scheme). These schemes have, for example, been applied to nonlinear solid mechanics in [211]. The extension to interface problems is possible by defining independent aggregated spaces on both sides of the interface [23]. In this work continuity of traces is weakly enforced (as Dirichlet data), which leads to robust schemes that can handle high-contrast problems.

The extension of aggregated finite elements to indefinite systems has been analyzed in [212]. We note that its application to stabilized finite element formulations that transform the indefinite system into a definite one is straightforward [213]. The stability of indefinite problems relies on so-called inf-sup conditions. These conditions are stringent and are only valid for very specific i.e., finite element spaces. For instance, the velocity and pressure finite element spaces in the (Navier–)Stokes problem have to match to satisfy this condition. Aggregation on these spaces does not generally preserve inf-sup stability. The aggregated mixed finite element method for the Stokes problem in [212] starts with the inf-sup stable pair $\mathcal{Q}_p(\mathcal{T}_h^{\text{in}}) \times \mathcal{P}_{p-1}^-(\mathcal{T}_h^{\text{in}})$, i.e., the velocity resides component-wise in $\mathcal{Q}_p(\mathcal{T}_h^{\text{in}}) \cap C^0(\mathcal{T}_h^{\text{in}})$ and the pressure is a discontinuous function in $\mathcal{P}_{p-1}^-(\mathcal{T}_h^{\text{in}})$. The extension defined above for continuous and discontinuous spaces leads to a pair that does not satisfy the inf-sup condition. Instead, a modified extension of the velocity components is proposed

in [212], which relies on serendipity finite elements. In some specific situations, additional pressure jump stabilization terms must also be added.

Aggregation has also been applied to enable stable explicit time stepping, such as in, e.g., [154] for the wave equation.

4.4 Ghost-Penalty Stabilization

In this section, we present common techniques for the weak stabilization of immersed finite element methods, often referred to as the *ghost penalty* and generally associated with the *Cut Finite Element Method* (CutFEM) [76, 77, 79]. In the ghost-penalty approach, additional terms are added to the bilinear form which provide control over the solution on the elements that intersect the boundary, while the original approximation space V_h is unmodified. These terms are denoted by $s_h(v_h, w_h)$ and provide a contribution to the operator norm in the form of $\|v_h\|_{s_h}^2 = s_h(v_h, v_h)$. The underlying idea is to control the variation of the approximation function across neighboring elements, such that the solution on small cut elements is controlled by the solution on interior elements. A standard way to achieve this is by controlling the jump in the (higher-order) normal derivatives across faces, by adding appropriately scaled stiffnesses to these jumps. Another approach is to add a volumetric penalization to the difference between a function itself and the extension of the function on a neighboring element. In the subsequent sections, both approaches are discussed.

4.4.1 Face-Based Stabilization

The most common stabilization term for cut elements is the face-based ghost penalty [76, 77], defined as

$$s_h(w_h, v_h) = \sum_{F \in \mathcal{F}_h^{\text{gh}}} \sum_{j=1}^p \tau_j h^{2j-1} (\llbracket \partial_n^j w_h \rrbracket, \llbracket \partial_n^j v_h \rrbracket)_F \quad (4.5)$$

where τ_j are positive constants, $\mathcal{F}_h^{\text{gh}}$ is the set of all interior faces that are shared by an element in $\mathcal{T}_h^{\text{cut}}$ (see Fig. 8), $(\cdot, \cdot)_F$ indicates the inner product over F , and

$$\llbracket \partial_n v_h \rrbracket = \partial_{n_1} v_{h,1} + \partial_{n_2} v_{h,2} \quad (4.6)$$

is the jump in the normal gradient across the face F . In this definition of the jump on face F , n_1 and $v_{h,1}$, and n_2 and $v_{h,2}$, correspond to the exterior unit normal and the function v_h on the elements T_1 and T_2 that share face F , respectively. Note that, therefore, $n_1 = -n_2$. With this penalty on the jumps in the normal derivatives between two neighboring elements, the norm of a function $v_h \in V_h$ on one of the elements can be controlled by the norm on the other. Accordingly, gradients




Fig. 8 Set of ghost-penalty faces $\mathcal{F}_h^{\text{gh}}$ marked in red. These faces are interior faces of the active mesh \mathcal{T}_h (yellow and blue elements) and belong to at least one cut element in $\mathcal{T}_h^{\text{cut}}$ (blue elements)




Fig. 9 Illustration of the extension and difference of $v_{h,1}$ and $v_{h,2}$ in (4.8)

on cut elements are controlled by gradients on full elements in the interior.

The parameters τ_j in the stabilization terms must be chosen in an appropriate way. With too small parameters, the effect is negligible, while with too large parameters, degrees of freedom can be constrained too strongly, potentially degenerating the accuracy of the approximation [214]. On Dirichlet boundaries, the required value to provide coercivity of the bilinear form correlates inversely with the Nitsche parameter. On Neumann boundaries, coercivity of the bilinear form does not depend on the stabilization term. Nevertheless, stabilization can still be applied to the Neumann boundary in order to repair the conditioning problems discussed in Sect. 3.1, for which a scaling with h^{2j+1} instead of h^{2j-1} in (4.5) suffices [215].

To elaborate the ghost-penalty approach, let us for simplicity consider the case of linear elements, i.e., $p = 1$. It then holds that

$$\begin{aligned} \|\nabla v_h\|_{T_1}^2 &\lesssim \|\nabla v_h\|_{T_2}^2 + h \|\llbracket \partial_n v_h \rrbracket\|_F^2, \\ \|v_h\|_{T_1}^2 &\lesssim \|v_h\|_{T_2}^2 + h^3 \|\llbracket \partial_n v_h \rrbracket\|_F^2 \quad v_h \in V_h \end{aligned} \quad (4.7)$$

To derive these estimates, we consider the element-wise restrictions $v_{h,i} = v_h|_{T_i}$ ($i \in \{1, 2\}$), both extended onto $T_1 \cup T_2$; see Fig. 9. The second estimate in (4.7) follows from

$$\begin{aligned} \|v_{h,1}\|_{T_1}^2 &\lesssim \|v_{h,1} - v_{h,2}\|_{T_1}^2 + \|v_{h,2}\|_{T_1}^2 \\ &\lesssim \|(\xi - \xi_F) \cdot \nabla(v_{h,1} - v_{h,2})\|_{T_1}^2 + \|v_{h,2}\|_{T_2}^2 \\ &\lesssim h^3 \|\partial_{n_i}(v_{h,1} - v_{h,2})\|_F^2 + \|v_{h,2}\|_{T_2}^2 \\ &= h^3 \|\llbracket \partial_n v \rrbracket\|_F^2 + \|v_{h,2}\|_{T_2}^2 \end{aligned} \quad (4.8)$$

with ∂_{n_i} indicating either of the normal derivatives. Here we first added and subtracted $v_{h,2}$ and used the triangle inequality. We then used the Taylor expansion $v_{h,1} - v_{h,2} = (\xi - \xi_F) \cdot \nabla(v_{h,1} - v_{h,2})$ at ξ with point $\xi_F \in F$ the projection of ξ along the normal direction, which holds since $v_{h,1} - v_{h,2} = 0$ on F by virtue of the continuity of v_h . Additionally, we used the estimate $\|v_{h,2}\|_{T_1} \lesssim \|v_{h,2}\|_{T_2}$ which holds since $v_{h,2}$ is polynomial. The first estimate follows in the same way, but is simpler since the gradients are piecewise constant, leading to an h scaling factor instead of the h^3 factor.

Using the pairwise bounds given in (4.7)—either directly or through a chain of neighboring elements—the degrees of freedom on cut elements in $\mathcal{T}_h^{\text{cut}}$ can be controlled by the degrees of freedom on interior elements $\mathcal{T}_h^{\text{in}}$ via the ghost penalty

$$\begin{aligned} \|\nabla v_h\|_{\Omega_h}^2 &\lesssim \|\nabla v_h\|_{\Omega}^2 + s_h(v_h, v_h), \\ \|v_h\|_{\Omega_h}^2 &\lesssim \|v_h\|_{\Omega}^2 + h^2 s_h(v_h, v_h) \quad v_h \in V_h \end{aligned} \quad (4.9)$$

To prove the first bound we proceed as

$$\begin{aligned} \|\nabla v_h\|_{\Omega_h}^2 &= \sum_{T \in \mathcal{T}_h} \|\nabla v_h\|_T^2 \\ &= \sum_{T \in \mathcal{T}_h^{\text{in}}} \|\nabla v_h\|_T^2 + \sum_{T \in \mathcal{T}_h^{\text{cut}}} \|\nabla v_h\|_T^2 \\ &\lesssim \sum_{T \in \mathcal{T}_h^{\text{in}}} \|\nabla v_h\|_T^2 + \left(\sum_{T \in \mathcal{T}_h^{\text{in}}} \|\nabla v_h\|_T^2 + s_h(v_h, v_h) \right) \\ &\lesssim \|\nabla v_h\|_{\Omega}^2 + s_h(v_h, v_h) \end{aligned} \quad (4.10)$$

The second bound in (4.9) can be established similarly. Note that the constants in these estimates will depend on the lengths of the paths connecting cut elements to interior elements, which is related to the properties and the resolution

of the boundary, as well as the properties of the mesh. For smooth boundaries and locally quasi-uniform meshes with sufficiently small mesh sizes, one can show that the length of these paths is uniformly bounded.

Conversely, the stabilization term can be bounded in terms of the L^2 -norm of the gradient of the approximation on the active domain Ω_h

$$s_h(v_h, v_h) \lesssim \|\nabla v_h\|_{\Omega_h} \quad v_h \in V_h \quad (4.11)$$

which follows from

$$\begin{aligned} h^{2j-1} \|[\partial_n^j v_h]\|_F^2 &\lesssim \sum_{i=1}^2 h^{2j-1} \|\partial_n^j v_{h,i}\|_F^2 \\ &\lesssim \sum_{i=1}^2 h^{2j-2} \|\nabla^j v_{h,i}\|_{T_i}^2 \\ &\lesssim \sum_{i=1}^2 \|\nabla v_{h,i}\|_{T_i}^2 = \|\nabla v_h\|_{T_1 \cup T_2}^2 \end{aligned} \quad (4.12)$$

where we first employed the triangle inequality, then a trace inverse estimate to pass from the face F to the neighboring elements T_1 and T_2 [100], and finally another inverse estimate to arrive at a bound in terms of ∇v_h [96]. Using the same inverse estimate it follows that

$$h^2 \|v_h\|_{s_h}^2 \lesssim h^2 \|\nabla v_h\|_{\Omega_h}^2 \lesssim \|v_h\|_{\Omega_h}^2 \quad (4.13)$$

The estimates (4.9), (4.11), and (4.13) are central to the coercivity and continuity of the (stabilized) bilinear form with a (global) Nitsche parameter that scales inversely with the mesh size of the *background* grid. The estimates are also essential for the condition number bounds derived in Sect. 4.5. The principal notion of the ghost-penalty method is encapsulated in the bounds in (4.9), which convey that the ghost penalty provides control over the solution on the entire active domain including the fictitious part. The bounds (4.11) and (4.13) in turn impart that the ghost-penalty term is suitably well behaved.

4.4.2 Element- and Patch-Based Stabilization Forms

An alternative to the face-based penalty is a penalty that is based on volumetric integrals. This is convenient for higher-order polynomials, since it avoids the computation of higher-order derivatives at element interfaces, which is a non-standard operation in many finite element codes. An overview and analysis of different types of such ghost-penalty formulations is presented in [216].

A stabilization term based on element integrals can be defined according to any of the two forms

$$\begin{aligned} s_{h,0}^{\text{elem}}(w_h, v_h) &= \sum_{F \in \mathcal{F}_h^{\text{gh}}} \tau_0 h^{-2} ([w_h], [v_h])_{T_h(F)}, \\ s_{h,1}^{\text{elem}}(w_h, v_h) &= \sum_{F \in \mathcal{F}_h^{\text{gh}}} \tau_1 ([\nabla w_h], [\nabla v_h])_{T_h(F)} \end{aligned} \quad (4.14)$$

where for each face $F \in \mathcal{F}_h^{\text{gh}}$, $T_h(F)$ is the pair of elements that share F and

$$[v_h] = v_{h,1} - v_{h,2}, \quad [\nabla v_h] = \nabla v_{h,1} - \nabla v_{h,2} \quad (4.15)$$

on $T \in T_h(F) = \{T_1, T_2\}$ with polynomials $v_{h,1}$ and $v_{h,2}$ extended in the canonical way; see Fig. 9. Note that, opposed to the existing volumetric integrals in the weak formulation introduced in Sect. 2.1, the integrals in the penalization are taken over the entire background elements including the fictitious parts and are not restricted to the physical domain Ω . The form in (4.14) is presented in [217, 218] and is obtained by example formulation 2 in [216, §2.4] in case the element pairs are chosen based on the faces in $\mathcal{F}_h^{\text{gh}}$.

We may also consider stabilization based on patches or macro elements, which is the form that was presented with the introduction of the ghost penalty in [79] and, for a certain choice of patches, is also considered in [218]. With this form of stabilization, each cut element is associated with a patch of neighboring elements that contains at least one element in the interior. We may then penalize the difference between the finite element function and a global polynomial on the patch. More precisely, let M_T be a patch of elements containing T , then we can define

$$\begin{aligned} s_{h,0}^{\text{patch}}(w_h, v_h) &= \sum_{T \in \mathcal{T}_h^{\text{cut}}} \tau_0 h^{-2} (w_h - P_p w_h, v_h - P_p v_h)_{M_T}, \\ s_{h,1}^{\text{patch}}(w_h, v_h) &= \sum_{T \in \mathcal{T}_h^{\text{cut}}} \tau_1 (\nabla(w_h - P_p w_h), \nabla(v_h - P_p v_h))_{M_T} \end{aligned} \quad (4.16)$$

where $P_p : V_h|_{M_T} \rightarrow \mathcal{P}_p(M_T)$ is the L^2 -projector or the Ritz projector, respectively. Note that by strictly following this formulation, patches containing multiple cut elements will be contained in the summation multiple times. In practice, every patch that contains cut elements is included only once, or the boundary is simply covered by a set of non-overlapping patches such as in [219].

Ghost-penalty stabilization can also be combined with the concept of aggregation to design weakly aggregated schemes, as proposed in [81, 216]. In these schemes, one makes use of the standard finite element space V_h on the active mesh, and instead of strongly imposing constraints via an extension operator of interior degrees of freedom, the difference between V_h and V_h^{ag} is weakly penalized. Recall the discrete extension operator $\mathcal{E}_h^{\text{ag}} : V_h^{\text{in}} \rightarrow V_h^{\text{ag}} \in V_h$ defined in Sect. 4.3 to define a

mapping $P_h^{\text{ag}} : V_h \ni v_h \mapsto \mathcal{E}_h^{\text{ag}}(v_h|_{\mathcal{T}_h^{\text{in}}}) \in V_h^{\text{ag}}$. The difference $v_h - P_h^{\text{ag}} v_h$ can then be penalized similarly as in (4.16)

$$\begin{aligned} s_h^{\text{ag}}(w_h, v_h) &= \sum_{T_h} \tau_0 h^{-2} (w_h - P_h^{\text{ag}} w_h, v_h - P_h^{\text{ag}} v_h)_T, \\ s_h^{\text{ag}}(w_h, v_h) &= \sum_{T_h} \tau_1 (\nabla(w_h - P_h^{\text{ag}} w_h), \nabla(v_h - P_h^{\text{ag}} v_h))_T \end{aligned} \quad (4.17)$$

These terms can be computed element-wise (in a similar manner as described in Sect. 4.3) and are similar to local-projection stabilization techniques; see, e.g., [220, 221]. We note that the (strong) aggregated finite element method of Sect. 4.3 is formally recovered in the limit $\tau_j \rightarrow \infty$. Thus, the stabilization method is also accurate for large values of τ_j . This last statement does not hold for other forms of the ghost penalty, in which a large parameter can constrain too many degrees of freedom, which inhibits the accuracy [214] and is referred to as locking in [216]. Detailed analyses and experimentation with this type of stabilization can be found in [81] and [216].

Stabilization techniques based on volumetric integrals have also been developed for multimesh discretizations [222–224] and for problems on manifolds [78]. Additionally, it is noteworthy that with the H^1 -product and with the projector $P_p = 0$ or $P_h^{\text{ag}} = 0$, the fictitious domain stiffness that is commonly employed in the finite cell method is recovered [225]. Fictitious domain stiffness requires $\tau_1 \ll 1$, however, as this stabilization term is not consistent with the original problem. Therefore, with a fictitious domain stiffness, stability and conditioning need to be balanced with consistency.

4.4.3 Implementation Aspects

To implement any of the aforementioned stabilization forms, one must be able to find face neighbors of given elements and assemble contributions from pairs of elements sharing a face. For higher-order elements, the face-based ghost penalty requires computation of higher-order derivatives. The element-based formulation may then be preferable. All the stabilization terms lead to additional coupling and, as a consequence, fill-in in the stiffness matrix [55, 226]. Systems stabilized by the ghost penalty have similar condition number bounds as boundary-fitted formulations, and therefore larger systems (i.e., finer grids) yield larger condition numbers and require more iterations in iterative solvers. This can be overcome by combining the ghost penalty with a multigrid solver; see the discussion on this in Sect. 3.2.

4.4.4 Extension to Other Problems

The ghost penalty can also be applied to stabilize time-dependent problems, for which generally the mass form needs to be stabilized as well. In that case the scaling is chosen in such a way that the stabilization terms scale with the mesh parameter in the same way as the mass matrix; see (4.13). Similar to Schwarz preconditioning and aggregation, the ghost penalty can also be applied to provide inf-sup stable discretization of mixed formulations such as that of the Stokes problem; see, e.g., [26, 32, 55, 227]. Furthermore, the ghost penalty can also be applied to stabilize discretizations of codimensional manifolds, such as partial differential equations posed on non-planar surfaces discretized with three-dimensional meshes; see, e.g., [78, 137, 228, 229].

4.5 A Unified Analysis of Aggregation and Stabilization

In this section, we present a framework for the analysis of stable higher-order immersed finite element methods. We consider aggregated and ghost-penalty-stabilized techniques at the same time, to identify the common parts and the main differences. It is recalled from Sect. 4.1 that when Nitsche's method is applied without any stability-enhancing technique, the bilinear form is only coercive with an *a priori* unbounded Nitsche parameter β , depending on the cut-element thickness. In that case, most of the properties and estimates presented in this section do not apply, and error bounds cannot be provided. Remark 4.2 at the end of this section discusses which specific analyses still hold and which no longer hold without any form of stabilization.

4.5.1 Methods and Formulations

We recall from (2.6) and (2.7) that the standard Nitsche finite element method takes the form

$$\begin{cases} \text{find } u_h \in V_h \text{ such that:} \\ a_h(u_h, v_h) = l_h(v_h) \quad \forall v_h \in V_h \end{cases} \quad (4.18)$$

with bilinear form $a_h(\cdot, \cdot)$ according to

$$\begin{aligned} a_h(w_h, v_h) &= (\nabla w_h, \nabla v_h)_\Omega - (w_h, \partial_n v_h)_{\partial\Omega_D} \\ &\quad - (\partial_n w_h, v_h)_{\partial\Omega_D} + (\beta w_h, v_h)_{\partial\Omega_D} \end{aligned} \quad (4.19)$$

We note that this is consistent with the strong formulation in (2.1), i.e., the exact solution u to the corresponding continuous problem satisfies

$$a_h(u, v_h) = l_h(v_h) \quad v_h \in V_h \quad (4.20)$$

provided that u is regular enough, for instance if $u \in H^2(\Omega)$. To accommodate both stabilized approaches presented in Sects. 4.3 and 4.4 in the same analysis, we consider the formulation

$$\begin{cases} \text{find } u_h \in V_{h,\star} \text{ such that:} \\ a_{h,\star}(u_h, v_h) = l_h(v_h) \quad \forall v_h \in V_{h,\star} \end{cases} \quad (4.21)$$

where the finite dimensional function space is defined by

$$V_{h,\star} = \begin{cases} V_h^{\text{ag}} \subset V_h & \text{aggregation} \\ V_h & \text{weak stabilization} \end{cases} \quad (4.22)$$

and the bilinear form is defined by

$$\begin{aligned} a_{h,\star}(w_h, v_h) = & a_\Omega(w_h, v_h) - (w_h, \partial_n v_h)_{\partial\Omega_D} \\ & - (\partial_n w_h, v_h)_{\partial\Omega_D} + (\beta w_h, v_h)_{\partial\Omega_D} \end{aligned} \quad (4.23)$$

with

$$\begin{aligned} a_\Omega(w_h, v_h) \\ = \begin{cases} (\nabla w_h, \nabla v_h)_\Omega & \text{aggregation} \\ (\nabla w_h, \nabla v_h)_\Omega + s_h(w_h, v_h) & \text{weak stabilization} \end{cases} \end{aligned} \quad (4.24)$$

that only requires β inversely proportional to the mesh size of the background grid. For the ghost-penalty method, the stabilization form is included in the definition of $a_\Omega(\cdot, \cdot)$ because the stabilization form $s_h(\cdot, \cdot)$ is naturally paired with the bulk or volumetric term $(\nabla \cdot, \nabla \cdot)_\Omega$ to enable bounds of the form (4.7) and (4.9).

4.5.2 Properties

For the analysis, we start from the following intrinsic properties of aggregation methods and ghost-penalty formulations:

- There is a constant independent of the mesh size and the cut elements (hidden in the binary operator \lesssim) such that

$$\|\nabla v_h\|_{\Omega_h}^2 \lesssim a_\Omega(v_h, v_h) = \|v_h\|_{a_\Omega}^2 \quad v_h \in V_{h,\star} \quad (4.25)$$

where as before $\Omega_h = \cup_{T \in \mathcal{T}_h} T$ is the union of all the active elements, such that $\Omega_h \supset \Omega$. See, respectively, (4.3) and (4.10) for these estimates in the context of aggregated and weakly stabilized methods. This property forms the basis of the effectiveness of weakly stabilized and aggregated methods. Essentially, it indicates that the solution on the extended domain Ω_h , including the fictitious parts outside Ω of active elements, is controlled by $a_\Omega(\cdot, \cdot)$. Consequently, this precludes the very small eigenvalues on small cut elements discussed in Sect. 3.1, and it provides control of the normal gradients which enables bounding the Nitsche parameter β from above

as discussed in Sect. 4.1 and 4.2. With aggregation this is achieved by directly precluding basis functions with a very small support in the physical domain as discussed in Sect. 4.3, while with the ghost penalty such functions are given a contribution in $a_\Omega(\cdot, \cdot)$ by $s_h(\cdot, \cdot)$ as discussed in Sect. 4.4.

- There is an interpolation operator $\pi_h : H^s(\Omega) \rightarrow V_{h,\star}$, $s \geq 2$, such that independent of how Ω intersects elements

$$\begin{aligned} \|\nabla^m(v - \pi_h v)\|_{\Omega_h} &\lesssim h^{s-m} \|v\|_{H^s(\Omega)} \\ 0 \leq m \leq s &\leq p+1 \end{aligned} \quad (4.26)$$

We use a continuous extension operator $\mathcal{E} : v \in H^s(\Omega) \rightarrow H^s(\mathbb{R}^d)$ to extend v outside of Ω . For ghost-penalty-stabilized methods, the interpolation operator is generally constructed by first extending the function v outside of Ω and then applying a standard interpolation operator, for instance the Cl  ment operator [230]. For the aggregated finite element space, a specific construction of the extension is used to establish the interpolation estimate [206, 231].

- The stabilization form $s_h(\cdot, \cdot)$ satisfies

$$\|\pi_h v\|_{s_h}^2 = s_h(\pi_h v, \pi_h v) \lesssim h^{s-1} \|v\|_{H^s(\Omega)}^2 \quad 1 \leq s \leq p+1 \quad (4.27)$$

independently of how elements are cut. This is a consistency assumption on the stabilization form $s_h(\cdot, \cdot)$, formulated in such a way that the stabilization form only acts directly on functions in the finite element space, and acts on other functions v indirectly via their projection $\pi_h v$. Note that the stabilization form $s_h(\cdot, \cdot)$ may for instance involve higher-order derivatives, which are not generally well defined for the exact solution or the solution to a dual problem used in an error analysis. This complication is avoided by formulating the consistency condition (4.27) in terms of projections onto the finite element space.

To provide an example we verify this consistency assumption (4.27) for the face-based stabilization form (4.5) in the case of linear elements. Consider a face $F \in \mathcal{F}_h$ shared by the two elements T_1 and T_2 . Then for any linear function $w \in \mathcal{P}_1(T_1 \cup T_2)$ we have

$$h \|\llbracket \partial_n \pi_h v \rrbracket\|_F^2 = h \|\llbracket \partial_n (\pi_h v - w) \rrbracket\|_F^2 \lesssim \sum_{i=1}^2 \|\nabla (\pi_h v - w)\|_{T_i}^2 \quad (4.28)$$

where we used a trace inverse estimate to pass from the face to the elements [100]. Using the Bramble-Hilbert lemma (see, e.g., [232]) we can choose w in such a way that

$$\begin{aligned} \sum_{i=1}^2 \|\nabla(\pi_h v - w)\|_{T_i}^2 &\leq \sum_{i=1}^2 \|\nabla(\pi_h v - v)\|_{T_i}^2 \\ &+ \|\nabla(v - w)\|_{T_i}^2 \lesssim \sum_{i=1}^2 h^2 \|v\|_{H^2(S(T_i))}^2 \end{aligned} \quad (4.29)$$

where $S(T_i)$ is the union of the set of elements in T_h that share a node with T .

Element aggregation and ghost-penalty stabilization are conceptually different in the manner in which they approach the conditions (4.25)–(4.27). Aggregation methods aim to satisfy condition (4.25) by restricting the approximation to a suitable subspace $V_h^{\text{ag}} \subset V_h$. On the other hand, condition (4.26) demands optimal approximation properties of the (sequence of) aggregated approximation spaces V_h^{ag} and, hence, the aggregated approximation spaces cannot be too small. Because the stabilization operator $s_h(\cdot, \cdot)$ vanishes in aggregation methods, condition (4.27) is trivially satisfied. Ghost-penalty stabilization requires that the stabilization term $s_h(\cdot, \cdot)$ is strong enough to comply with condition (4.25). On the other hand, condition (4.27) insists that the stabilization cannot be too strong, as this will affect the solution. Because ghost-penalty stabilization imposes no auxiliary conditions on the approximation spaces, condition (4.26) is trivially satisfied.

4.5.3 Coercivity and Continuity

We define the following energy norms

$$\|v\|_h^2 = \|\nabla v\|_{\Omega}^2 + \|h^{-\frac{1}{2}} v\|_{\partial\Omega_D}^2 + \|h^{\frac{1}{2}} \partial_n v\|_{\partial\Omega_D}^2 \quad (4.30)$$

$$\|v_h\|_{h,\star}^2 = a_{\Omega}(v_h, v_h) + \|h^{-\frac{1}{2}} v_h\|_{\partial\Omega_D}^2 + \|h^{\frac{1}{2}} \partial_n v_h\|_{\partial\Omega_D}^2 \quad (4.31)$$

and note that

$$\|v_h\|_{h,\star}^2 = \begin{cases} \|v_h\|_h^2 & \text{aggregation} \\ \|v_h\|_h^2 + \|v_h\|_{s_h}^2 & \text{weak stabilization} \end{cases} \quad (4.32)$$

Note that these norms are different from the β -norm that was applied in [1] and presented in (4.1) in Sect. 4.1. First, $\|\cdot\|_{h,\star}$ employs h^{-1} , which is uniformly bounded, instead of β , which is not bounded in case no stabilization technique is applied. Second, in case of ghost-penalty stabilization, $\|\cdot\|_{h,\star}$ also includes a ghost-penalty contribution. Consequently, whereas the β -norm admits a natural analysis of unstabilized forms, the $\|v_h\|_{h,\star}$ -norm naturally accompanies weakly stabilized and aggregated formulations.

Lemma 4.1 *The form $a_{h,\star}(\cdot, \cdot)$ is continuous on $V_{h,\star} \oplus H^2(\Omega)$*

$$a_{h,\star}(w, v) \leq \|w\|_{h,\star} \|v\|_{h,\star} \quad w, v \in V_{h,\star} \oplus H^2(\Omega) \quad (4.33)$$

and, if β is chosen large enough, $a_{h,\star}(\cdot, \cdot)$ is coercive on $V_{h,\star}$ independently of how elements are intersected by the boundary

$$\|v_h\|_{h,\star}^2 \lesssim a_{h,\star}(v_h, v_h) \quad v_h \in V_{h,\star} \quad (4.34)$$

Proof Under the notion that with aggregation or weak stabilization $\beta \lesssim h^{-1}$, the continuity follows directly from the Cauchy-Schwarz inequality together with the observation that $\partial_n v \in L^2(\partial\Omega_D)$ for $v \in V_{h,\star} \oplus H^2(\Omega)$ and therefore $\|\partial_n v\|_{\partial\Omega_D}$ is well defined. For the coercivity we use the element-wise trace inverse inequality [100] (note that the right hand side considers the full background element T and not only the intersection with Ω)

$$h \|\partial_n v\|_{T \cap \partial\Omega_D}^2 \lesssim \|\nabla v\|_T^2 \quad v \in \mathcal{P}_p(T) \quad (4.35)$$

and (4.25) to conclude that independent of how the elements are cut

$$\begin{aligned} \|h^{\frac{1}{2}} \partial_n v_h\|_{\partial\Omega_D}^2 &\lesssim \sum_{T \in \mathcal{T}_h^{\text{cut}}} \|\nabla v_h\|_T^2 \lesssim \sum_{T \in \mathcal{T}_h} \|\nabla v_h\|_T^2 \\ &\lesssim \|\nabla v_h\|_{\Omega_h}^2 \lesssim \|v_h\|_{a_{\Omega}}^2 \quad v_h \in V_{h,\star} \end{aligned} \quad (4.36)$$

where $\mathcal{T}_h^{\text{cut}}$ is the set of elements that intersect the boundary. Next, for any $\delta > 0$, we have

$$\begin{aligned} 2(\partial_n v, v)_{\partial\Omega_D} &\leq 2\|\partial_n v\|_{\partial\Omega_D} \|v\|_{\partial\Omega_D} \\ &\leq \|(h\delta)^{\frac{1}{2}} \partial_n v\|_{\partial\Omega_D}^2 + \|(h\delta)^{-\frac{1}{2}} v\|_{\partial\Omega_D}^2 \end{aligned} \quad (4.37)$$

which leads to

$$\begin{aligned} a_{h,\star}(v_h, v_h) &= \|v_h\|_{a_{\Omega}}^2 - 2(\partial_n v_h, v_h)_{\partial\Omega_D} + \|\beta^{\frac{1}{2}} v_h\|_{\partial\Omega_D}^2 \\ &\geq \|v_h\|_{a_{\Omega}}^2 - \|(h\delta)^{\frac{1}{2}} \partial_n v_h\|_{\partial\Omega_D}^2 \\ &\quad - \|(h\delta)^{-\frac{1}{2}} v_h\|_{\partial\Omega_D}^2 + \|\beta^{\frac{1}{2}} v_h\|_{\partial\Omega_D}^2 \\ &\geq (1 - \delta C) \|v_h\|_{a_{\Omega}}^2 + \|(\beta - (h\delta)^{-1})^{\frac{1}{2}} v_h\|_{\partial\Omega_D}^2 \\ &\quad v_h \in V_{h,\star} \end{aligned} \quad (4.38)$$

where C is the hidden constant between the first and final expressions in (4.36) that does not depend on how elements intersect Ω . Taking δ small enough to guarantee that $\delta C \lesssim 1/2$ and β large enough such that $\beta - (h\delta)^{-1} > h^{-1}/2$ on each element in $\mathcal{T}_h^{\text{cut}}$, the coercivity follows. \square

Assuming that $l_h(v_h)$ is continuous, i.e., $l_h(v_h) \lesssim \|v_h\|_{h,\star}$ for $v_h \in V_{h,\star}$, we may now apply the Lax–Milgram lemma to show that there is a unique and stable solution $u_h \in V_{h,\star}$ to (4.21).

4.5.4 Error Estimates

Conditions (4.25)–(4.27) ensure that the corresponding immersed finite-element approximation according to (4.18) has optimal approximation properties in the energy norm.

Theorem 4.1 *Consider the immersed finite-element approximation (4.18). Assume that conditions (4.25)–(4.27) hold. The approximate solution u_h satisfies the energy-norm error estimate independently of how elements are cut*

$$\|u - u_h\|_h \lesssim h^p \|u\|_{H^{p+1}(\Omega)} \quad (4.39)$$

Proof We start by splitting the error in two parts using the interpolation operator

$$\|u - u_h\|_h \lesssim \|u - \pi_h u\|_h + \|\pi_h u - u_h\|_h \quad (4.40)$$

We first consider the first term and show that

$$\|u - \pi_h u\|_h \lesssim h^p \|u\|_{H^{p+1}(\Omega)} \quad (4.41)$$

To establish (4.41), we note that by virtue of a standard trace inequality (see, e.g., [232], note that again the full background element T is considered)

$$\|v\|_{T \cap \partial \Omega_D}^2 \lesssim h^{-1} \|v\|_T^2 + h \|\nabla v\|_T^2 \quad v \in H^1(T) \quad (4.42)$$

it holds that

$$\begin{aligned} \|v\|_h^2 &= \|\nabla v\|_\Omega^2 + \|h^{-\frac{1}{2}} v\|_{\partial \Omega_D}^2 + \|h^{\frac{1}{2}} \partial_n v\|_{\partial \Omega_D}^2 \\ &\lesssim \|\nabla v\|_{\Omega_h}^2 + \|h^{-1} v\|_{\Omega_h}^2 + \|h \nabla^2 v\|_{\Omega_h}^2 \end{aligned} \quad (4.43)$$

for $v \in V_{h,\star} + H^2(\Omega)$. Setting $v = u - \pi_h u$ and using (4.26), we obtain (4.41). Equation (4.43) applies the definition of $\|\cdot\|_h$ given in (4.30).

Considering the second term in (4.40), we note that the definition of $\|\cdot\|_{h,\star}$ and the coercivity property (4.34) imply

$$\begin{aligned} \|\pi_h u - u_h\|_h &\lesssim \|\pi_h u - u_h\|_{h,\star} \\ &\lesssim \sup_{w_h \in V_{h,\star} \setminus \{0\}} \frac{a_{h,\star}(\pi_h u - u_h, w_h)}{\|w_h\|_{h,\star}} \end{aligned} \quad (4.44)$$

In the case of the aggregation method, we have $a_{h,\star}(\cdot, \cdot) = a_h(\cdot, \cdot)$ and $\|\cdot\|_{h,\star} = \|\cdot\|_h$ and it immediately follows that independent of how elements intersect Ω

$$\begin{aligned} a_{h,\star}(\pi_h u - u_h, w_h) &= a_h(\pi_h u - u_h, w_h) \\ &= a_h(\pi_h u, w_h) - l_h(w_h) \\ &= a_h(\pi_h u - u, w_h) \\ &\lesssim \|\pi_h u - u\|_h \|w_h\|_h \\ &\lesssim h^p \|u\|_{H^{p+1}(\Omega)} \|w_h\|_h \end{aligned} \quad (4.45)$$

where we used, consecutively, the consistency relation (4.20) to replace $l_h(w_h)$ by $a_h(u, w_h)$, continuity of a_h according to (4.33), and the interpolation error estimate (4.41). In the case of ghost-penalty stabilization, it holds that independent of how elements are cut

$$\begin{aligned} a_{h,\star}(\pi_h u - u_h, w_h) &= a_{h,\star}(\pi_h u, w_h) - l_h(w_h) \\ &= a_h(\pi_h u, w_h) - l_h(w_h) + s_h(\pi_h u, w_h) \\ &= a_h(\pi_h u - u, w_h) + s_h(\pi_h u, w_h) \\ &\lesssim \|\pi_h u - u\|_h \|w_h\|_h + \|\pi_h u\|_{s_h} \|w_h\|_{s_h} \\ &\lesssim \left(\|\pi_h u - u\|_h^2 + \|\pi_h u\|_{s_h}^2 \right)^{1/2} \|w_h\|_{h,\star} \\ &\lesssim h^p \|u\|_{H^{p+1}(\Omega)} \|w_h\|_{h,\star} \end{aligned} \quad (4.46)$$

where we used the interpolation error estimate (4.41) and the weak consistency condition (4.27) on $s_h(\cdot, \cdot)$. From (4.44)–(4.46) it follows that the second term in (4.40) is also bounded in accordance with (4.39). \square

4.5.5 Condition Number Estimate

Under certain (non-restrictive) conditions, both aggregation and ghost-penalty stabilization engender stiffness matrices that exhibit similar condition numbers as conventional boundary-fitted finite element methods. To analyze the condition numbers corresponding to both methods, we consider the basis $\{\phi_i\}_{i=1}^N$ in $V_{h,\star}$, so that any $v_h \in V_{h,\star}$ can be expanded as

$$v_h = \sum_{i=1}^N y_i \phi_i \quad (4.47)$$

where $\mathbf{y} \in \mathbb{R}^N$ is the coefficient vector. It is to be noted that with ghost-penalty stabilization the basis $V_{h,\star} = V_h$ is a standard nodal basis, while with aggregation $V_{h,\star} = V_h^{\text{ag}}$ is the basis that is modified at the boundary in such a way that the support of each basis function has a sufficiently large intersection with Ω , more precisely $|\text{supp}(\phi_i) \cap \Omega| \sim h^d$. The stiffness matrix $\mathbf{A} \in \mathbb{R}^{N \times N}$ is defined as

$$\mathbf{z}^T \mathbf{A} \mathbf{y} = a_{h,\star}(w_h, v_h) \quad w_h, v_h \in V_{h,\star} \quad (4.48)$$

and for the considered symmetric bilinear forms, the condition number of the stiffness matrix \mathbf{A} corresponds to

$$\kappa(\mathbf{A}) = \frac{\lambda_{\max}}{\lambda_{\min}} \quad (4.49)$$

where λ_{\max} and λ_{\min} are the maximum and minimum eigenvalues of \mathbf{A} , respectively. To estimate the condition number we start from the following basic bounds, which hold by construction for both aggregated and weakly stabilized methods:

- There is a constant (hidden in the binary relation \lesssim) such that independent of how elements are intersected by the boundary

$$\|v_h\|_{h,\star} \lesssim h^{-1} \|v_h\|_{\Omega_h} \quad v_h \in V_{h,\star} \quad (4.50)$$

Both for weakly stabilized and aggregated methods, this bound follows from standard inverse inequalities; see (4.11) for the inverse estimate of the ghost-penalty form.

- There is a (hidden) constant independent of how elements are cut such that

$$\|v_h\|_{\Omega_h} \lesssim \|v_h\|_{h,\star} \quad v_h \in V_{h,\star} \quad (4.51)$$

This is a Poincaré inequality, requiring aggregation of the approximation space or ghost-penalty stabilization of the bilinear form to gain control on the approximate solution outside Ω .

- There are constants such that the following equivalence holds:

$$h^d \|\mathbf{y}\|_2^2 \sim \|v_h\|_{\Omega_h}^2 \quad v_h \in V_{h,\star} \quad (4.52)$$

For weakly stabilized methods we have $V_{h,\star} = V_h$ and this is a standard estimate that follows from the quasi-uniformity of the mesh and the sufficient linear independence of the basis functions. For the aggregated finite element space $V_{h,\star} = V_h^{\text{ag}}$, this estimate holds by virtue of the construction of the extension operator.

Theorem 4.2 Consider the stiffness matrix in (4.48). Assume that the bilinear form $a_{h,\star}(\cdot, \cdot)$ is coercive according to (4.34) and that conditions (4.50)–(4.52) hold. The condition number $\kappa(\mathbf{A})$ of the stiffness matrix satisfies the following bound

$$\kappa(\mathbf{A}) \lesssim h^{-2} \quad (4.53)$$

independently of how elements intersect Ω .

Proof To estimate the condition number according to (4.49), we use the following Rayleigh quotient characterization of the eigenvalues to bound the maximum (resp. minimum) eigenvalue from above (resp. below):

$$\lambda_{\max} = \max_{\mathbf{y} \in \mathbb{R}^N \setminus \{0\}} \frac{\mathbf{y}^T \mathbf{A} \mathbf{y}}{\|\mathbf{y}\|_2^2}, \quad \lambda_{\min} = \min_{\mathbf{y} \in \mathbb{R}^N \setminus \{0\}} \frac{\mathbf{y}^T \mathbf{A} \mathbf{y}}{\|\mathbf{y}\|_2^2} \quad (4.54)$$

We bound the Rayleigh quotient from above as follows

$$\frac{\mathbf{y}^T \mathbf{A} \mathbf{y}}{\|\mathbf{y}\|_2^2} \sim \frac{a_{h,\star}(v_h, v_h)}{h^{-d} \|v_h\|_{\Omega_h}^2} = h^d \frac{a_{h,\star}(v_h, v_h)}{\|v_h\|_{h,\star}^2} \frac{\|v_h\|_{h,\star}^2}{\|v_h\|_{\Omega_h}^2} \lesssim h^{d-2} \quad (4.55)$$

where we used, consecutively, the definition (4.48) of the stiffness matrix, the equivalence (4.52), and continuity

of $a_{h,\star}(\cdot, \cdot)$ according to (4.33) and the inverse inequality (4.50). Conversely, to bound the quotient from below, we infer from the equivalence relation (4.52), coercivity condition (4.34), and the Poincaré inequality (4.51) that

$$\frac{\mathbf{y}^T \mathbf{A} \mathbf{y}}{\|\mathbf{y}\|_2^2} \gtrsim \frac{\|v_h\|_{h,\star}^2}{h^{-d} \|v_h\|_{\Omega_h}^2} \gtrsim h^d \quad (4.56)$$

Combining the upper and lower bounds, we obtain the desired estimate

$$\kappa = \frac{\lambda_{\max}}{\lambda_{\min}} \lesssim \frac{h^{d-2}}{h^d} \lesssim h^{-2} \quad (4.57)$$

□

Remark 4.2 Comparison to Nitsche's method without stabilization

We here briefly discuss which aspects of Sect. 4.5 do, and do not, apply when Nitsche's method is applied in an immersed setting without an additional stability-enhancing technique. First and foremost, the intrinsic property (4.25) will not be satisfied. Consequently, the normal gradients on the boundary are not sufficiently controlled by the gradients in the bulk, such that coercivity only holds with locally very large (i.e., *a priori* unbounded) values for the Nitsche parameter β . Nevertheless, with the Nitsche parameter large enough, in the β -norm both (4.33) and (4.34) hold (i.e., respectively continuity and coercivity), and only the continuous dependence on the data in the Lax–Milgram lemma can formally not be verified as the right-hand-side operator $l_h(\cdot)$ is not bounded in the β -norm such that $\|u_h\|_\beta$ cannot be *a priori* bounded. The error estimate provided in (4.39) does not hold for unstabilized formulations in combination with the β -norm, as $\beta \sim h_{T_\Omega}^{-1} \gtrsim h^{-1}$ such that (4.43), and consequently (4.41), do not apply. Note that (4.41) is also employed in the proof of the bound on the second part of the right-hand-side of (4.40). As discussed in Sect. 4.1, this conclusion is also obtained in [1], which shows that a best approximation property in the β -norm can be demonstrated following Céa's lemma, but that the problem lies in the error estimates because a bound on the error of the best approximation in the β -norm cannot be provided. As already demonstrated in Sect. 3.1, condition number bounds cannot be provided and tailored preconditioning is required. Related to the analysis presented in this section, both (4.50) and (4.51) are contingent on ghost-penalty stabilization or aggregation. For (4.50) the bound cannot be provided as there is no *a priori* upper bound on β . Nevertheless, this lack of an upper bound to the eigenvalues is generally not the cause of the conditioning problems of these methods, and this type of large eigenvalues only occurs on highly irregularly shaped cut elements. The main problem with the conditioning is that without additional stabilization, (4.51) does not hold, and as

demonstrated in Sect. 3.1 it is this lack of a lower bound to the eigenvalues that generally causes the ill-conditioning of unstabilized immersed methods.

5 Discussion and Conclusion

In this review, we considered the problems resulting from small cut elements in immersed finite element methods, as well as the solutions that have been developed to counter these problems. This section assesses the current state of the field and presents a discussion about future research directions. Similarities and differences between the different techniques that are discussed in this manuscript are treated in the respective sections, and we do not provide a qualitative comparison here. This is because, first, the presented techniques pursue different goals related to the different problems on small cut elements, i.e., stability and conditioning. Second, the approaches are not mutually exclusive and combinations can be made. In Sect. 4 the combination of aggregation and ghost-penalty stabilization is already discussed, and similarly it is possible to combine it with, e.g., the minimal stabilization procedure, which fixes the stability but not the conditioning, with a preconditioning approach.

The field of immersed finite element methods has substantially grown and developed over the past decade. The employed formulations have matured, and are currently based on rigorous theoretical and mathematical foundations. Furthermore, efficient parallel implementations have been developed that enable the application to large problems into the billions of degrees of freedom. Certain aspects are still challenging, however, which particularly manifest in the application to real-world problems. These challenging aspects, and the corresponding unresolved questions, form the basis of future research directions that we believe are important for the further advancement of the field of immersed finite element methods. In this section we discuss these open research questions in a general sense, and indicate it if these are specifically related to one of the techniques that are discussed in this manuscript.

While nearly all publications on immersed methods provide qualitative arguments to compare immersed and boundary-fitted discretizations (such as in the first paragraph of the introduction of this manuscript, a detailed discussion is included in [233]), quantitative comparisons are scarce. This is largely caused by the difficulty of such a comparison, which requires the assessment of aspects that are difficult to correlate such as level of automation, computational cost, parallelizability, accuracy, and robustness. Furthermore, comparisons are heavily test-case specific, particularly since the cost and the accuracy of boundary-fitted methods strongly depend on the complexity and quality of mesh-generation, and additionally both approaches are generally not (equally well) implemented in a single code. Most

quantitative comparisons only compare the accuracy versus the number of degrees of freedom in benchmark computations, such as [11, 106, 234] and the benchmark treated in both [55, 226], and do not show large differences. A comprehensive, unbiased, and rigorous study to quantitatively compare these techniques on all relevant aspects would therefore be very welcome to the field, where it should nevertheless be noted that such a study will never be conclusive since contributions continue to be made.

An aspect that requires further consideration to advance the field of immersed finite element methods, is the detection of small features in the geometry or the solution. In boundary-fitted discretizations, the mesh size is commonly connected to the resolution of geometrical features, and defeaturing may be required to simplify the meshing procedure [235]. The decoupling of the geometry description from the finite element mesh provides immersed formulations with flexibility in this regard, and makes the consideration of adaptive discretizations natural. It is notable that grid refinements in the vicinity of domain boundaries are commonly considered in the *hp*-adaptive finite cell method [21, 106, 236, 237]. Such refinements are not required along all boundaries, but mainly in regions where small solution features such as large gradients are expected. For efficiency purposes, we therefore consider the further advancement of error estimators and error-estimation-based adaptivity schemes tailored to immersed formulations as presented in [238, 239] a desired development. One complication in regard of *a posteriori* error estimates is that the bilinear form is only coercive on discrete spaces, and that the dependence of the bilinear form on the discretization (both on the mesh size and on the order) complicates saturation assumptions. Besides the accuracy considerations related to enriched approximation spaces along small geometrical features, stabilized methods also require a certain level of geometry resolution, which makes local adaptive schemes well suited for the application of these techniques to problems on (scan-based) porous domains, as the refinements reduce the ratio of elements that are cut. Since element aggregation strongly ties degrees of freedom on badly cut elements to well posed degrees of freedom, local refinement procedures can be beneficial in cases where a large portion (or in certain regions even the majority) of the active mesh consists of badly cut elements. This is similar for ghost-penalty stabilization terms on domains with (relative to the mesh) thin structures, which can potentially make the domain artificially stiff in such cases. Note that in these cases the choice of the parameters η^* (with aggregation; see Sect. 4.3) and τ (with ghost-penalty stabilization; see Sect. 4.4) is particularly critical, and that in many cases problems can be avoided by a adequate parameter choice. Nevertheless, the accurate approximation of problems with very thin features in an unfitted setting remains a difficult problem for any

method, and deserves further attention. Schwarz preconditioners are not affected by the matter of mesh resolution, but it should be noted that this only pertains to the solving of the linear system. With a natural boundary condition on the unfitted boundary and a preconditioner to resolve the conditioning problems, a relatively low level of mesh resolution can therefore lead to reasonable solutions. It should be mentioned that this can still result in artificial couplings as described in e.g., [51, 172], however. With Nitsche's method to enforce essential boundary conditions on the unfitted boundary, such as with flow problems in porous media, the application of a preconditioner without an additional stabilization technique can lead to inaccurate approximations, as described in Sect. 4.

While this manuscript focuses on single-field problems, mixed formulations, in particular flow problems, are an important application area of immersed finite element methods. The techniques presented in this manuscript are all applicable to such problems, and can retain or precondition inf-sup stable formulations (see, e.g., [212] for aggregation, [26, 240] for ghost penalty, and [64] for preconditioning). An unresolved topic is the application of (pointwise) divergence-conforming or geometry-preserving discretizations for incompressible problems. Compatible pairs of approximation spaces that result in pointwise divergence-free approximations for boundary-fitted discretizations, generally lose this property in an unfitted setting. Furthermore, stabilization techniques complicate the formulation of pointwise divergence-free methods. Ghost-penalty terms interfere with the enforcement of incompressibility in the bilinear form, and it has not been thoroughly investigated how the existing discrete aggregation and extension operators affect the compatibility of structure-preserving pairs of approximation spaces. Besides these comments about $H(\text{div})$ -conforming discretizations for the conservation of geometry, similar considerations can be made about $H(\text{curl})$ -conforming discretizations when solving the Maxwell equations in electromagnetism.

An aspect that relates to both of the previous topics of mesh adaptivity and flow problems is the consideration of anisotropic grid refinements. To accurately capture boundary layers at a reasonable computational cost, boundary layer meshes are commonly anisotropic with a much finer resolution normal to the boundary. While the aspect of boundary layers in immersed methods has already been considered in multiple studies (see, e.g., [28, 241]), present understanding of this aspect is incomplete, and this is a topic that requires further investigation.

While only briefly considered in Sect. 2.2.3, a topic that is strongly related to small cut elements is explicit time integration for problems in dynamics. This is particularly interesting for crash simulations, as such problems are commonly solved by explicit time steppers [159]. Although tools exist

that result in stable time steps that are not affected by the cut elements (i.e., element aggregation, ghost penalty terms in the mass matrix, or lumping of the mass matrix [154, 158, 159]), the research on this topic is not exhaustive. In particular, a comprehensive analysis of stable time-step sizes with different choices for the mass and stiffness forms (i.e., boundary condition enforcement technique, ghost-penalty stabilization of both forms, lumping of the mass matrix, etc.) does, to the authors' knowledge, not yet exist.

Another open question deals with the choice, or the optimization, of parameters in the presented approaches. This does not only concern, e.g., the parameters η^* and τ that are explicitly contained in element aggregation and the ghost penalty, but also hidden parameters (or choices) in the setup such as the choice of the index blocks with Schwarz preconditioning. While these parameters generally do not affect the asymptotic behavior, well-chosen (or, conversely, poorly-chosen) parameters can significantly affect the performance pre-asymptotically. We therefore consider an optimal choice of (hidden) parameters an important topic for future research. This specifically applies to the dependence of this optimality on the basis-function type and the discretization order and, in particular for discretizations with local refinements or local order elevations, to the option of choosing parameters locally or element-wise (similar to the Nitsche parameter without stabilization).

Over the past years, many well-developed, well-documented, and user-friendly open source codes have become available that aid the application of immersed finite element methods, a comprehensive list of which is beyond the scope of this work. These toolboxes support all facets of immersed finite element computations, such as pre-processing steps (in particular integration and the application of boundary conditions), the solution process itself, and also post-processing, where often a subtriangulation is employed like this is also done for higher-order boundary-fitted finite elements. Something that is not yet available is a graphical user interface that makes immersed finite element methods accessible also to non-expert users. Besides the developments in open source codes, also commercial software has adopted (aspects of) immersed methods. To provide some examples: in Ansys Fluent the Immersed Boundary Method (IBM) is available [66, 242]; Ansys LS-DYNA facilitates the application of trimmed isogeometric analysis in both shell and solid models for static and dynamic analyses [243, 244]; in Abaqus it is possible to (partially) include and exclude elements from a computation with *element progressive activation* [245, 246], which is mainly intended for the simulation of additive manufacturing; and the platform Hyperganic [247] employs immersed finite element techniques to perform simulations. The implementation of the stabilization methods discussed in this work can require tools that are not generally available in existing software and codes. For

instance, face-based ghost-penalty stabilization requires the computation of jumps in normal gradients up to the order of approximation. For quadratic and higher-order discretizations, this functionality is generally not available, even for software frameworks that include the possibility to perform face loops and compute jumps for standard Discontinuous Galerkin (DG) formulations. The implementation of aggregation or patch-based ghost-penalty stabilization requires the computation of aggregate-wise terms. These terms can be implemented using standard element-wise terms if the finite element software provides functionality to enforce linear constraints. Many codes already provide this functionality as this is essential for local grid refinements and periodic boundary conditions (e.g., deal.ii [248], FEniCS [249], Gridap [250], FEMPAR [209], Netgen [251], or Dune [252]). While (geometric) multigrid routines are intrusive to the finite element software, the implementation of tailored Schwarz preconditioners is exceptionally non-intrusive and only requires the matrix itself and the basis function connectivity. Furthermore, as additive Schwarz preconditioning is a well-established concept also in boundary-fitted finite element methods, these preconditioners are available in many codes. The construction of the index blocks is not generally the same in existing software, however, and it should be noted that the effectivity for immersed methods relies on the specific construction of these blocks. Additionally, advanced procedures in which e.g., index blocks are only devised for basis functions with support on cut elements, require the availability of this information or a routine to extract it.

By virtue of the stabilization, aggregation, and preconditioning techniques as discussed in this review, immersed finite element methods have developed into reliable and versatile computational analysis instruments. In our assessment, the potential of immersed methods and, in particular, their ability to provide new simulation workflows with minimal user intervention, has not been fully exploited yet. Another recommendation therefore pertains to the exploration of immersed methods in new application fields, such as hyperbolic systems, diffuse interface models, poro-elasticity, computational aeroacoustics, and phase field fracture.

Declarations

Conflict of interest On behalf of all authors, the corresponding author states that there is no conflict of interest.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in

the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. de Prenter F, Lehrenfeld C, Massing A (2018) A note on the stability parameter in Nitsche's method for unfitted boundary value problems. *Comput Math Appl* 75(12):4322–4336
2. Düster A, Parvizián J, Yang Z, Rank E (2008) The finite cell method for three-dimensional problems of solid mechanics. *Comput Methods Appl Mech Eng* 197(45–48):3768–3782
3. Schillinger D, Dedè L, Scott MA, Evans JA, Borden MJ, Rank E, Hughes TJR (2012) An isogeometric design-through-analysis methodology based on adaptive hierarchical refinement of NURBS, immersed boundary methods, and T-spline CAD surfaces. *Comput Methods Appl Mech Eng* 249–252:116–150
4. Schillinger D, Ruess M, Zander N, Bazilevs Y, Düster A, Rank E (2012) Small and large deformation analysis with the p- and B-spline versions of the Finite Cell Method. *Comput Mech* 50:445–478
5. Schillinger D, Düster A, Rank E (2012) The hp-d-adaptive finite cell method for geometrically nonlinear problems of solid mechanics. *Int J Numer Meth Eng* 89(9):1171–1202
6. Zander N, Kollmannsberger S, Ruess M, Yosibash Z, Rank E (2012) The Finite Cell Method for linear thermoelasticity. *Comput Math Appl* 64(11):3527–3541
7. Jouliaian M, Duczek S, Gabbert U, Düster A (2014) Finite and spectral cell method for wave propagation in heterogeneous materials. *Comput Mech* 54:661–675
8. Ruess M, Schillinger D, Özcan AI, Rank E (2014) Weak coupling for isogeometric analysis of non-matching and trimmed multipatch geometries. *Comput Methods Appl Mech Eng* 269:46–71
9. Rank E, Ruess M, Kollmannsberger S, Schillinger D, Düster A (2012) Geometric modeling, isogeometric analysis and the finite cell method. *Comput Methods Appl Mech Eng* 249–252:104–115
10. Ruess M, Schillinger D, Bazilevs Y, Varduhn V, Rank E (2013) Weakly enforced essential boundary conditions for NURBS-embedded and trimmed NURBS geometries on the basis of the finite cell method. *Int J Numer Meth Eng* 95(10):811–846
11. Rank E, Kollmannsberger S, Sorger C, Düster A (2011) Shell Finite Cell Method: a high order fictitious domain approach for thin-walled structures. *Comput Methods Appl Mech Eng* 200(45–46):3200–3209
12. Schmidt R, Wüchner R, Bletzinger K-U (2012) Isogeometric analysis of trimmed NURBS geometries. *Comput Methods Appl Mech Eng* 241–244:93–111
13. Guo Y, Ruess M (2015) Weak Dirichlet boundary conditions for trimmed thin isogeometric shells. *Computers & Mathematics with Applications* 70(7):1425–1440
14. Bauer AM, Breitenberger M, Philipp B, Wüchner R, Bletzinger K-U (2017) Embedded structural entities in NURBS-based isogeometric analysis. *Comput Methods Appl Mech Eng* 325:198–218
15. Guo Y, Ruess M, Schillinger D (2017) A parameter-free variational coupling approach for trimmed isogeometric thin shells. *Comput Mech* 59:693–715
16. Dolbow JE, Harari I (2009) An efficient finite element method for embedded interface problems. *Int J Numer Meth Eng* 78(2):229–252
17. Schillinger D, Rank E (2011) An unfitted hp-adaptive finite element method based on hierarchical B-splines for interface problems of complex geometry. *Comput Methods Appl Mech Eng* 200(47–48):3358–3380

18. Annavarapu C, Hautefeuille M, Dolbow JE (2012) A robust Nitsche's formulation for interface problems. *Comput Methods Appl Mech Eng* 225–228:44–54
19. Wadbro E, Zahedi S, Kreiss G, Berggren M (2013) A uniformly well-conditioned, unfitted Nitsche method for interface problems. *BIT Numer Math* 53:791–820
20. Hansbo P, Larson MG, Zahedi S (2014) A cut finite element method for a Stokes interface problem. *Appl Numer Math* 85:90–114
21. Elhaddad M, Zander N, Bog T, Kudela L, Kollmannsberger S, Kirschke J, Baum T, Ruess M, Rank E (2017) Multi-level hp-finite cell method for embedded interface problems with application in biomechanics. *International Journal for Numerical Methods in Biomedical Engineering* 34(4):e2951
22. Wu H, Xiao Y (2019) An unfitted hp-interface penalty finite element method for elliptic interface problems. *J Comput Math* 37(3):316–339
23. Neiva E, Badia S (2021) Robust and scalable h-adaptive aggregated unfitted finite elements for interface elliptic problems. *Comput Methods Appl Mech Eng* 380:113769
24. Chen Z, Li K, Xiang X (2021) An adaptive high-order unfitted finite element method for elliptic interface problems. *Numer Math* 149:507–548
25. Burman E, Elfverson D, Hansbo P, Larson MG, Larsson K (2019) Hybridized CutFEM for elliptic interface problems. *SIAM J Sci Comput* 41(5):A3354–A3380
26. Massing A, Larson MG, Logg A, Rognes ME (2014) A stabilized Nitsche fictitious domain method for the Stokes problem. *J Sci Comput* 61:604–628
27. Schott B, Wall WA (2014) A new face-oriented stabilized XFEM approach for 2D and 3D incompressible Navier-Stokes equations. *Comput Methods Appl Mech Eng* 276:233–265
28. Xu F, Schillinger D, Kamensky D, Varduhn V, Wang C, Hsu M-C (2016) The tetrahedral finite cell method for fluids: Immersogeometric analysis of turbulent flow around complex geometries. *Comput Fluids* 141:135–154
29. Schott B, Rasthofer U, Gravemeier V, Wall WA (2015) A face-oriented stabilized Nitsche-type extended variational multiscale method for incompressible two-phase flow. *Int J Numer Meth Eng* 104(7):721–748
30. Schott B, Shahmiri S, Kruse R, Wall WA (2016) A stabilized Nitsche-type extended embedding mesh approach for 3D low- and high-Reynolds-number flows. *Int J Numer Meth Fluids* 82(6):289–315
31. Hsu M-C, Wang C, Xu F, Herrema AJ, Krishnamurthy A (2016) Direct immersogeometric fluid flow analysis using B-rep CAD models. *Computer Aided Geometric Design* 43:143–158
32. Hoang T, Verhoosel CV, Auricchio F, van Brummelen EH, Reali A (2017) Mixed Isogeometric Finite Cell Methods for the Stokes problem. *Comput Methods Appl Mech Eng* 316:400–423
33. Massing A, Schott B, Wall WA (2018) A stabilized Nitsche cut finite element method for the Oseen problem. *Comput Methods Appl Mech Eng* 328:262–300
34. Winter M, Schott B, Massing A, Wall WA (2018) A Nitsche cut finite element method for the Oseen problem with general Navier boundary conditions. *Comput Methods Appl Mech Eng* 330:220–252
35. Bazilevs Y, Hsu M-C, Scott MA (2012) Isogeometric fluid-structure interaction analysis with emphasis on non-matching discretizations, and with application to wind turbines. *Comput Methods Appl Mech Eng* 249:28–41
36. Burman E, Fernández MA (2014) An unfitted Nitsche method for incompressible fluid-structure interaction using overlapping meshes. *Comput Methods Appl Mech Eng* 279:497–514
37. Rüberg T, Cirak F (2014) A fixed-grid b-spline finite element technique for fluid-structure interaction. *Int J Numer Meth Fluids* 74(9):623–660
38. Massing A, Larson MG, Logg A, Rognes ME (2015) A Nitsche-based cut finite element method for a fluid-structure interaction problem. *Communications in Applied Mathematics and Computational Science* 10(2):97–120
39. Kadapa C, Dettmer WG, Perić D (2016) A fictitious domain/distributed Lagrange multiplier based fluid-structure interaction scheme with hierarchical B-spline grids. *Comput Methods Appl Mech Eng* 301:1–27
40. Kadapa C, Dettmer WG, Perić D (2017) A stabilised immersed boundary method on hierarchical b-spline grids for fluid-rigid body interaction with solid-solid contact. *Comput Methods Appl Mech Eng* 318:242–269
41. Wang Y, Jimack PK, Walkley MA (2017) A one-field monolithic fictitious domain method for fluid-structure interactions. *Comput Methods Appl Mech Eng* 317:1146–1168
42. Wu MCH, Kamensky D, Wang C, Herrema AJ, Xu F, Pigazzini MS, Verma A, Marsden AL, Bazilevs Y, Hsu M-C (2017) Optimizing fluid-structure interaction systems with immersogeometric analysis and surrogate modeling: application to a hydraulic arresting gear. *Comput Methods Appl Mech Eng* 316:668–693
43. Kadapa C, Dettmer WG, Perić D (2018) A stabilised immersed framework on hierarchical b-spline grids for fluid-flexible structure interaction with solid-solid contact. *Comput Methods Appl Mech Eng* 335:472–489
44. Hsu M-C, Kamensky D, Bazilevs Y, Sacks MS, Hughes TJR (2014) Fluid-structure interaction analysis of bioprosthetic heart valves: significance of arterial wall deformation. *Comput Mech* 54:1055–1071
45. Kamensky D, Hsu M-C, Schillinger D, Evans JA, Aggarwal A, Bazilevs Y, Sacks MS, Hughes TJR (2015) An immersogeometric variational framework for fluid-structure interaction: application to bioprosthetic heart valves. *Comput Methods Appl Mech Eng* 284:1005–1053
46. Hsu M-C, Kamensky D, Xu F, Kiendl J, Wang C, Wu MCH, Mineroff J, Reali A, Bazilevs Y, Sacks MS (2015) Dynamic and fluid-structure interaction simulations of bioprosthetic heart valves using parametric design with T-splines and Fung-type material models. *Comput Mech* 55:1211–1225
47. Kamensky D, Hsu M-C, Yu Y, Evans JA, Sacks MS, Hughes TJR (2017) Immersogeometric cardiovascular fluid-structure interaction analysis with divergence-conforming B-splines. *Comput Methods Appl Mech Eng* 314:408–472
48. Yang Z, Kollmannsberger S, Düster A, Ruess M, Garcia EG, Burgkart R, Rank E (2011) Non-standard bone simulation: interactive numerical analysis by computational steering. *Comput Vis Sci* 14:207–216
49. Ruess M, Tal D, Trabelsi N, Yosibash Z, Rank E (2012) The finite cell method for bone simulations: verification and validation. *Biomech Model Mechanobiol* 11:425–437
50. Düster A, Sehlhorst H-G, Rank E (2012) Numerical homogenization of heterogeneous and cellular materials utilizing the finite cell method. *Comput Mech* 50:413–431
51. Verhoosel CV, Van Zwieten GJ, Van Rietbergen B, De Borst R (2015) Image-based goal-oriented adaptive isogeometric analysis with application to the micro-mechanical modeling of trabecular bone. *Comput Methods Appl Mech Eng* 284:138–164
52. Varduhn V, Hsu M-C, Ruess M, Schillinger D (2016) The tetrahedral finite cell method: higher-order immersogeometric analysis on adaptive non-boundary-fitted meshes. *Int J Numer Meth Eng* 107(12):1054–1079
53. Duczek S, Berger H, Gabbert U (2015) The Finite Pore Method: a new approach to evaluate gas pores in cast parts by combining

- computed tomography and the finite cell method. *Int J Cast Met Res* 28(4):221–228
54. Würkner M, Ducek S, Berger H, Köppe H, Gabbert U (2018) A software platform for the analysis of porous die-cast parts using the finite cell method. Springer, pp 327–341
 55. Hoang T, Verhoosel CVCV, Qin C-Z, Auricchio F, Reali A, van Brummelen EH (2019) Skeleton-stabilized immersogeometric analysis for incompressible viscous flow problems. *Comput Methods Appl Mech Eng* 344:421–450
 56. Parvizian J, Düster A, Rank E (2012) Topology optimization using the finite cell method. *Optim Eng* 13:57–78
 57. Nadal E, Ródenas JJ, Albelda J, Tur M, Tarancón JE, Fuenmayor FJ (2013) Efficient finite element methodology based on Cartesian grids: application to structural shape optimization. *Abstract Appl Anal* p 953786
 58. van Dijk NP, Maute K, Langelaar M, van Keulen F (2013) Level-set methods for structural topology optimization: a review. *Struct Multidiscip Optim* 48:437–472
 59. Bandara K, Rüberg T, Cirak F (2016) Shape optimisation with multiresolution subdivision surfaces and immersed finite elements. *Comput Methods Appl Mech Eng* 300:510–539
 60. Groen JP, Langelaar M, Sigmund O, Ruess M (2017) Higher-order multi-resolution topology optimization using the finite cell method. *Int J Numer Meth Eng* 110(10):903–920
 61. Villanueva CH, Maute K (2017) CutFEM topology optimization of 3D laminar incompressible flow problems. *Comput Methods Appl Mech Eng* 320:444–473
 62. Burman E, Hansbo P, Larson MG (2018) A cut finite element method with boundary value correction. *Math Comput* 87:633–657
 63. Jomo JN, de Prenter F, Elhaddad M, D'Angella D, Verhoosel CV, Kollmannsberger S, Kirschke JS, Nübel V, van Brummelen EH, Rank E (2019) Robust and parallel scalable iterative solutions for large-scale finite cell analyses. *Finite Elem Anal Des* 163:14–30, Figures reproduced with permission from Elsevier
 64. de Prenter F, Verhoosel CV, van Brummelen EH (2019) Preconditioning immersed isogeometric finite element methods with application to flow problems. *Computer Methods Appl Mech Eng* 348:604–631, Figures reproduced with permission from Elsevier
 65. Peskin CS (1972) Flow patterns around heart valves: a numerical method. *J Comput Phys* 10(2):252–271
 66. Mittal R, Iaccarino G (2005) Immersed boundary methods. *Annu Rev Fluid Mech* 37:239–261
 67. Melenk JM, Babuška I (1996) The partition of unity finite element method: basic theory and applications. *Comput Methods Appl Mech Eng* 139(1–4):289–314
 68. Babuška I, Melenk JM (1997) The partition of unity method. *Int J Numer Meth Eng* 40(4):727–758
 69. Duarte CA, Babuška I, Oden JT (2000) Generalized finite element methods for three-dimensional structural mechanics problems. *Comput Struct* 77(2):215–232
 70. Strouboulis T, Babuška I, Copps K (2000) The design and analysis of the Generalized Finite Element Method. *Comput Methods Appl Mech Eng* 181(1–3):43–69
 71. Belytschko T, Black T (1999) Elastic crack growth in finite elements with minimal remeshing. *Int J Numer Meth Eng* 45(5):601–620
 72. Moës N, Dolbow JE, Belytschko T (1999) A finite element method for crack growth without remeshing. *Int J Numer Meth Eng* 46(1):131–150
 73. Hansbo A, Hansbo P (2002) An unfitted finite element method, based on Nitsche's method, for elliptic interface problems. *Comput Methods Appl Mech Eng* 191(47–48):5537–5552
 74. Parvizian J, Düster A, Rank E (2007) Finite cell method. *Comput Mech* 41:121–133
 75. Schillinger D, Ruess M (2015) The finite cell method: a review in the context of higher-order structural analysis of CAD and image-based geometric models. *Arch Comput Methods Eng* 22:391–455
 76. Burman E, Hansbo P (2012) Fictitious domain finite element methods using cut elements: II. A stabilized Nitsche method. *Appl Numer Math* 62(4):328–341
 77. Burman E, Claus S, Hansbo P, Larson MG, Massing A (2015) CutFEM: discretizing geometry and partial differential equations. *Int J Numer Meth Eng* 104(7):472–501
 78. Burman E, Hansbo P, Larson MG, Massing A (2018) Cut finite element methods for partial differential equations on embedded manifolds of arbitrary codimensions. *ESAIM* 52(6):2247–2282
 79. Burman E (2010) Ghost penalty. *CR Math* 348(21–22):1217–1220
 80. Badia S, Verdugo F, Martín AF (2018) The aggregated unfitted finite element method for elliptic problems. *Comput Methods Appl Mech Eng* 336:533–553
 81. Badia S, Neiva E, Verdugo F (2022) Linking ghost penalty and aggregated unfitted methods. *Comput Methods Appl Mech Eng* 388:114232
 82. Marco O, Sevilla R, Zhang Y, Ródenas JJ, Tur M (2015) Exact 3D boundary representation in finite element analysis based on cartesian grids independent of the geometry. *Int J Numer Meth Eng* 103(6):445–468
 83. Höllig K, Reif U, Wipperf J (2001) Weighted extended B-spline approximation of Dirichlet problems. *SIAM J Numer Anal* 39(2):442–462
 84. Höllig K, Apprich C, Streit A (2005) Introduction to the Web-method and its applications. *Adv Comput Math* 23:215–237
 85. Sanches RAK, Bornemann PB, Cirak F (2011) Immersed b-spline (i-spline) finite element method for geometrically complex domains. *Comput Methods Appl Mech Eng* 200(13–16):1432–1445
 86. Johansson A, Larson MG (2013) A high order discontinuous Galerkin Nitsche method for elliptic problems with fictitious boundary. *Numer Math* 123:607–628
 87. Hughes TJR, Cottrell JA, Bazilevs Y (2005) Isogeometric analysis: CAD, finite elements, NURBS, exact geometry and mesh refinement. *Comput Methods Appl Mech Eng* 194(39–41):4135–4195
 88. Marussig B, Hughes TJR (2018) A review of trimming in isogeometric analysis: challenges, data exchange and simulation aspects. *Arch Comput Methods Eng* 25:1059–1127
 89. Main A, Scovazzi G (2018) The shifted boundary method for embedded domain computations. Part II: linear advection-diffusion and incompressible Navier-Stokes equations. *J Comput Phys* 372:996–1026
 90. Main A, Scovazzi G (2018) The shifted boundary method for embedded domain computations. Part I: Poisson and Stokes problems. *J Comput Phys* 372:972–995
 91. Atallah NM, Canuto C, Scovazzi G (2020) The second-generation Shifted Boundary Method and its numerical analysis. *Comput Methods Appl Mech Eng* 372:113341
 92. Atallah NM, Canuto C, Scovazzi G (2021) Analysis of the shifted boundary method for the Poisson problem in domains with corners. *Math Comput* 90:2041–2069
 93. Bramble JH, Dupont T, Thomée V (1972) Projection methods for Dirichlet's problem in approximating polygonal domains with boundary-value corrections. *Math Comput* 26:869–879
 94. Burman E, Ern A (2018) An unfitted hybrid high-order method for elliptic interface problems. *SIAM J Numer Anal* 56(3):1525–1546
 95. Badia S, Droniou J, Yemm L (2022) Conditioning of a hybrid high-order scheme on meshes with small faces. *J Sci Comput* 92(71)

96. Ern A, Guermond J-L (2004) Theory and practice of finite elements. Springer, Berlin
97. Evans LC (2010) Partial differential equations. American Mathematical Society
98. Nitsche J (1971) Über ein Variationsprinzip zur Lösung von Dirichlet-Problemen bei Verwendung von Teilräumen, die keinen Randbedingungen unterworfen sind. In Abhandlungen aus dem mathematischen Seminar der Universität Hamburg, pp 9–15
99. Embar A, Dolbow JE, Harari I (2010) Imposing Dirichlet boundary conditions with Nitsche's method and spline based finite elements. *Int J Numer Meth Eng* 83(7):877–898
100. Warburton T, Hesthaven JS (2003) On the constants in hp-finite element trace inverse inequalities. *Comput Methods Appl Mech Eng* 192(25):2765–2773
101. Evans JA, Hughes TJR (2013) Explicit trace inequalities for isogeometric analysis and parametric hexahedral finite elements. *Numer Math* 123:259–290
102. Cottrell JA, Hughes TJR, Bazilevs Y (2009) Isogeometric analysis: toward integration of CAD and FEA. Wiley, New York
103. Korshunova N, Jomo J, Léko G, Reznik D, Balázs P, Kollmannsberger S (2020) Image-based material characterization of complex microarchitected additively manufactured structures. *Comput Math Appl* 80(11):2462–2480
104. Abedian A, Parvizián J, Düster A, Khademyzadeh H, Rank E (2013) Performance of different integration schemes in facing discontinuities in the Finite Cell Method. *Int J Comput Methods* 10(3):1350002
105. Fries T-P, Belytschko T (2010) The extended/generalized finite element method: an overview of the method and its applications. *Int J Numer Meth Eng* 84(3):253–304
106. Düster A, Rank E, Szabó B (2017) The p-version of the finite element and finite cell methods. Wiley, New York, pp 1–35
107. Müller B, Kummer F, Oberlack M, Wang Y (2012) Simple multi-dimensional integration of discontinuous functions with application to level set methods. *Int J Numer Meth Eng* 92(7):637–651
108. Stavrev A, Nguyen LH, Shen R, Varduhn V, Behr M, Elgeti S, Schilling D (2016) Geometrically accurate, efficient, and flexible quadrature techniques for the tetrahedral finite cell method. *Comput Methods Appl Mech Eng* 310:646–673
109. Divi SC, Verhoosel CV, Auricchio F, Realí A, van Brummelen EH (2020) Error-estimate-based adaptive integration for immersed isogeometric analysis. *Computers & Mathematics with Applications* 80(11):2481–2516
110. Fries T-P, Omerović S (2016) Higher-order accurate integration of implicit geometries. *Int J Numer Meth Eng* 106(5):323–371
111. Fries T-P, Omerović S, Schöllhammer D, Steidl J (2017) Higher-order meshing of implicit geometries—Part I: integration and interpolation in cut elements. *Comput Methods Appl Mech Eng* 313:759–784
112. Omerović S, Fries T-P (2017) Conformal higher-order remeshing schemes for implicitly defined interface problems. *Int J Numer Meth Eng* 109(6):763–789
113. Fries T-P, Omerović S, Schöllhammer D, Stanford JW (2018) A fictitious domain method with higher-order accurate integration in cut elements. In: Proceedings of the 6th European conference on computational mechanics (ECCM 6) and the 7th European Conference on Computational Fluid Dynamics (ECFD 7), pp 3994–3405
114. Sevilla R, Fernández-méndez S, Huerta A (2008) NURBS-enhanced finite element method (NEFEM). *Int J Numer Meth Eng* 76(1):56–83
115. Sevilla R, Fernández-méndez S, Huerta A (2011) NURBS-enhanced finite element method (NEFEM). *Arch Comput Methods Eng* 18(441):441–484
116. Lehrenfeld C (2016) High order unfitted finite element methods on level set domains using isoparametric mappings. *Comput Methods Appl Mech Eng* 300:716–733
117. Kudela L, Zander N, Bog T, Kollmannsberger S, Rank E (2015) Efficient and accurate numerical quadrature for immersed boundary methods. *Adv Model Simul Eng Sci* 2(10)
118. Kudela L, Zander N, Kollmannsberger S, Rank E (2016) Smart octrees: Accurately integrating discontinuous functions in 3D. *Comput Methods Appl Mech Eng* 306:406–426
119. Hubrich S, Di Stolfo P, Kudela L, Kollmannsberger S, Rank E, Schröder A, Düster A (2017) Numerical integration of discontinuous functions: moment fitting and smart octree. *Comput Mech* 60:863–881
120. Mousavi SE, Sukumar N (2011) Numerical integration of polynomials and discontinuous functions on irregular convex polygons and polyhedrons. *Comput Mech* 47:535–554
121. Müller B, Kummer F, Oberlack M (2013) Highly accurate surface and volume integration on implicit domains by means of moment-fitting. *Int J Numer Meth Eng* 96(8):512–528
122. Müller B, Krämer-Eis S, Kummer F, Oberlack M (2017) A high-order discontinuous Galerkin method for compressible flows with immersed boundaries. *Int J Numer Meth Eng* 110(1):3–30
123. Joulaian M, Hubrich S, Düster A (2016) Numerical integration of discontinuities on arbitrary domains based on moment fitting. *Comput Mech* 57:979–999
124. Hubrich S, Düster A (2019) Numerical integration for nonlinear problems of the finite cell method using an adaptive scheme based on moment fitting. *Comput Math Appl* 77(7):1983–1997
125. Düster A, Hubrich S (2020) Adaptive integration of cut finite elements and cells for nonlinear structural analysis, pp 31–73. Springer, New York
126. Longva A, Löschner F, Kugelstadt T, Fernández-Fernández JA, Bender J (2020) Higher-order finite elements for embedded simulation. *ACM Trans Graph* 39(6)
127. Ventura G (2006) On the elimination of quadrature subcells for discontinuous functions in the eXtended Finite-Element Method. *Int J Numer Meth Eng* 66(5):761–795
128. Abedian A, Düster A (2019) Equivalent Legendre polynomials: numerical integration of discontinuous functions in the finite element methods. *Comput Methods Appl Mech Eng* 343:690–720
129. Ventura G, Gracie R, Belytschko T (2009) Fast integration and weight function blending in the extended finite element method. *Int J Numer Meth Eng* 77(1):1–29
130. Jonsson T, Larson MG, Larsson K (2017) Cut finite element methods for elliptic problems on multipatch parametric surfaces. *Comput Methods Appl Mech Eng* 324:366–394
131. Saye RI (2015) High-order quadrature methods for implicitly defined surfaces and volumes in hyperrectangles. *SIAM J Sci Comput* 37(2):A993–A1019
132. Cui T, Leng W, Liu H, Zhang L, Zheng W (2020) High-order numerical quadratures in a tetrahedron with an implicitly defined curved interface. *ACM Trans Math Softw* 46(1):3
133. Chin EB, Sukumar N (2020) An efficient method to integrate polynomials over polytopes and curved solids. *Computer Aided Geometric Design* 82:101914
134. Strang G, Fix GJ (1973) An analysis of the finite element method. Prentice-Hall, Hoboken
135. Strang G (1972) Variational crimes in the finite element method. Academic Press, Cambridge, pp 689–710
136. Burman E, Hansbo P, Larson MG, Zahedi S (2016) Cut finite element methods for coupled bulk-surface problems. *Numer Math* 133:203–231
137. Burman E, Hansbo P, Larson MG, Massing A (2017) A cut discontinuous Galerkin method for the Laplace-Beltrami operator. *IMA J Numer Anal* 37(1):138–169

138. Babuška I (1973) The finite element method with penalty. *Math Comput* 27:221–228
139. van Miegroet L, Duysinx P (2007) Stress concentration minimization of 2D filets using X-FEM and level set description. *Struct Multidiscip Optim* 33:425–438
140. Jiang W, Annavarapu C, Dolbow JE, Harari I (2015) A robust Nitsche's formulation for interface problems with spline-based finite elements. *Int J Numer Meth Eng* 104(7):676–696
141. Fernández-Méndez S, Huerta A (2004) Imposing essential boundary conditions in mesh-free methods. *Comput Methods Appl Mech Eng* 193(12–14):1257–1275
142. Burman E (2012) A penalty-free nonsymmetric Nitsche-type method for the weak imposition of boundary conditions. *SIAM J Numer Anal* 50(4):1959–1981
143. Boiveau T, Burman E (2016) A penalty-free Nitsche method for the weak imposition of boundary conditions in compressible and incompressible elasticity. *IMA J Numer Anal* 36(2):770–795
144. Schillinger D, Harari I, Hsu M-C, Kamensky D, Stoter SKF, Yu Y, Zhao Y (2016) The non-symmetric Nitsche method for the parameter-free imposition of weak boundary and coupling conditions in immersed finite elements. *Comput Methods Appl Mech Eng* 309:625–652
145. Codina R, Baiges J (2009) Approximate imposition of boundary conditions in immersed boundary methods. *Int J Numer Meth Eng* 80(11):1379–1405
146. Babuška I (1973) The finite element method with Lagrangian multipliers. *Numer Math* 20:179–192
147. Bertsekas DP (1982) *Constrained optimization and Lagrange multiplier methods*. Athena Scientific, Nashua
148. Burman E, Hansbo P (2010) Fictitious domain finite element methods using cut elements: I. A stabilized Lagrange multiplier method. *Comput Methods Appl Mech Eng* 199(41–44):2680–2686
149. Baiges J, Codina R, Henke F, Shahmiri S, Wall WA (2012) A symmetric method for weakly imposing Dirichlet boundary conditions in embedded finite element meshes. *Int J Numer Meth Eng* 90(5):636–658
150. Tur M, Albelda J, Nadal E, Ródenas JJ (2014) Imposing Dirichlet boundary conditions in hierarchical Cartesian meshes by means of stabilized Lagrange multipliers. *Int J Numer Meth Eng* 98(6):399–417
151. Kollmannsberger S, Özcan A, Baiges J, Ruess M, Rank E, Reali A (2015) Parameter-free, weak imposition of Dirichlet boundary conditions and coupling of trimmed and non-conforming patches. *Int J Numer Meth Eng* 101(9):670–699
152. de Prenter F, Verhoosel CV, van Zwieten GJ, van Brummelen EH (2017) Condition number analysis and preconditioning for the finite cell method. *Comput Methods Appl Mech Eng* 316:297–327
153. Huang P, Wu H, Xiao Y (2017) An unfitted interface penalty finite element method for elliptic interface problems. *Comput Methods Appl Mech Eng* 323:439–460
154. Burman E, Hansbo P, Larson MG (2022) Explicit time stepping for the wave equation using CutFEM with discrete extension. *SIAM J Sci Comput* 44(3):A1254–A1289
155. Harari I, Albocher U (2018) Spectral investigations of Nitsche's method. *J Sci Comput* 81:1472–1492
156. Harari I, Albocher U (2019) Complementary solutions of Nitsche's method. *Finite Elem Anal Des* 145:20–31
157. Albocher U, Harari I (2021) Spectral aspects of Nitsche's method on nonconforming meshes. *Mech Res Commun* 112:103611
158. Leidinger LF, Breitenberger M, Bauer AM, Hartmann S, Wüchner R, Bletzinger K-U, Duddeck F, Song L (2019) Explicit dynamic isogeometric B-Rep analysis of penalty-coupled trimmed NURBS shells. *Comput Methods Appl Mech Eng* 351:891–927
159. Leidinger LF (2020) Explicit isogeometric B-Rep analysis for nonlinear dynamic crash simulations. PhD thesis, Technical University of Munich
160. Golub GH, van Loan CF (2013) *Matrix computations*. John Hopkins University Press, Baltimore
161. Saad Y (2003) *Iterative methods for sparse linear systems*. SIAM
162. Greenbaum A (1997) *Iterative methods for solving linear systems*. SIAM
163. Barrett R, Berry MW, Chan TF, Demmel J, Donato J, Dongarra J, Eijkhout V, Pozo R, Romine C, van der Vorst H (1994) *Templates for the solution of linear systems: building blocks for iterative methods*. SIAM
164. Béchet E, Minnebo H, Moës N, Burgardt B (2005) Improved implementation and robustness study of the X-FEM for stress analysis around cracks. *Int J Numer Meth Eng* 64(8):1033–1056
165. Menk A, Bordas SPA (2011) A robust preconditioning technique for the extended finite element method. *Int J Numer Meth Eng* 85(13):1609–1632
166. Hiriyur B, Tuminaro RS, Waisman H, Boman EG, Keyes DE (2012) A quasi-algebraic multigrid approach to fracture problems based on extended finite elements. *SIAM J Sci Comput* 34(2):A603–A626
167. Berger-Vergiat L, Waisman H, Hiriyur B, Tuminaro RS, Keyes D (2012) Inexact Schwarz-algebraic multigrid preconditioners for crack problems modeled by extended finite element methods. *Int J Numer Meth Eng* 90(3):311–328
168. Waisman H, Berger-Vergiat L (2013) An adaptive domain decomposition preconditioner for crack propagation problems modeled by XFEM. *Int J Multiscale Comput Eng* 11(6):633–654
169. Lang C, Makhija D, Doostan K, Maute A (2014) A simple and efficient preconditioning scheme for Heaviside enriched XFEM. *Comput Mech* 54:1357–1374
170. Lehrenfeld C, Reusken A (2017) Optimal preconditioners for Nitsche-XFEM discretizations of interface problems. *Numer Math* 135:313–332
171. Ludescher T, Gross S, Reusken A (2020) A multigrid method for unfitted finite element discretizations of elliptic interface problems. *SIAM J Sci Comput* 42(1):A318–A342
172. de Prenter F, Verhoosel CV, van Brummelen EH, Evans JA, Messe C, Benzaken JD, Maute K (2020) Multigrid solvers for immersed finite element methods and immersed isogeometric analysis. *Comput Mech* 65:807–838
173. Saberi S, Vogel A, Meschke G (2020) Parallel finite cell method with adaptive geometric multigrid. In: *Euro-Par 2020: parallel processing*, pp 578–593
174. Jomo J, Oztoprak O, de Prenter F, Zander N, Kollmannsberger S, Rank E (2021) Hierarchical multigrid approaches for the finite cell method on uniform and multi-level hp-refined grids. *Comput Methods Appl Mech Eng* 386:114075
175. Badia S, Verdugo F (2018) Robust and scalable domain decomposition solvers for unfitted finite element methods. *J Comput Appl Math* 344:740–759
176. Badia S, Martín AF, Principe J (2016) Multilevel balancing domain decomposition at extreme scales. *SIAM J Sci Comput* 38(1):C22–C52
177. Kothari H, Krause R (2021) A multigrid method for a Nitsche-based extended finite element method. *Int J Comput Vis Sci Eng*
178. Nüssing A (2018) Fitted and unfitted finite element methods for solving the EEG forward problem. PhD thesis, University of Münster
179. Kothari H (2020) Multilevel solution strategies for unfitted finite element methods. PhD thesis, Università della Svizzera Italiana

180. Kothari H, Krause R (2021) Multigrid and saddle-point preconditioners for unfitted finite element modelling of inclusions. In: Proceedings of the 14th World congress on computational mechanics (WCCM) European community on computational methods in applied sciences (ECCOMAS) congress 2020
181. Kothari H, Krause R (2022) A generalized multigrid method for solving contact problems in Lagrange multiplier based unfitted Finite Element Method. *Comput Methods Appl Mech Eng* 392:114630
182. Gross S, Reusken A (2021) Optimal preconditioners for a Nitsche stabilized fictitious domain finite element method. *arXiv:2107.01182 [math.NA]*
183. Gross S, Reusken A (2022) Analysis of optimal preconditioners for CutFEM. *arXiv:2202.09069 [math.NA]*
184. de Dios B Ayuso, Dunn K, Sarkis M, Scacchi S (2020) Dirichlet–Neumann preconditioning for stabilised unfitted discretization of high contrast problems. In: Domain decomposition methods in science and engineering XXV, pp 67–74
185. Matsokin AM, Nepomnyashchikh SV (1985) The Schwarz alternation method in a subspace. *Russian Mathematics (Izvestiya Vysshikh Uchebnykh Zavedenii. Matematika)* 29(10):78–84
186. Lions PL (1988) On the Schwarz alternating method. I. In: Proceedings of the first international symposium on domain decomposition methods for partial differential equations, pp 1–42
187. Smith BF, Bjørstad PE, Gropp WD (1996) Domain decomposition: parallel multilevel methods for elliptic partial differential equations. Cambridge University Press, Cambridge
188. Toselli A, Widlund OB (2005) Domain decomposition methods: algorithms and theory. Springer, Berlin
189. Dauge M, Düster A, Rank E (2015) Theoretical and Numerical Investigation of the Finite Cell Method. *J Sci Comput* 65:1039–1064
190. Buffa A, Puppi R, Vázquez R (2020) A minimal stabilization procedure for isogeometric methods on trimmed geometries. *SIAM J Numer Anal* 58(5):2711–2735
191. Haslinger J, Renard Y (2009) A new fictitious domain approach inspired by the extended finite element method. *SIAM J Numer Anal* 47(2):1474–1499
192. Antolin P, Buffa A, Puppi R, Wei X (2021) Overlapping multi-patch isogeometric method with minimal stabilization. *SIAM J Sci Comput* 43(1):A330–A354
193. Puppi R (2021) Isogeometric discretizations of the Stokes problem on trimmed geometries. *arXiv:2012.15582 [math.NA]*
194. Tur M, Albelda J, Marco O, Ródenas JJ (2015) Stabilized method of imposing Dirichlet boundary conditions using a recovered stress field. *Comput Methods Appl Mech Eng* 296:352–375
195. Navarro-Jiménez JM, Nadal E, Tur M, Martínez-Casas J, Ródenas JJ (2020) On the use of stabilization techniques in the Cartesian grid finite element method framework for iterative solvers. *Int J Numer Meth Eng* 121(13):3004–3020
196. Elfverson D, Larson MG, Larsson K (2019) A new least squares stabilized Nitsche method for cut isogeometric analysis. *Comput Methods Appl Mech Eng* 349:1–16
197. Larsson K, Kollmannsberger S, Rank E, Larson MG (2022) The finite cell method with least squares stabilized Nitsche boundary conditions. *Comput Methods Appl Mech Eng* 393:114792
198. Rüberg T, Cirak F (2012) Subdivision-stabilised immersed b-spline finite elements for moving boundary flows. *Comput Methods Appl Mech Eng* 209–212:266–283
199. Rüberg T, Cirak F, García-Aznar JM (2016) An unstructured immersed finite element method for nonlinear solid mechanics. *Advanced Modeling and Simulation in Engineering Sciences* 3(1):22
200. Marussig B, Zechner J, Beer G, Fries T-P (2017) Stable isogeometric analysis of trimmed geometries. *Comput Methods Appl Mech Eng* 316:497–521
201. Marussig B, Hiemstra R, Hughes TJR (2018) Improved conditioning of isogeometric analysis matrices for trimmed geometries. *Comput Methods Appl Mech Eng* 334:79–110
202. Burman E, Hansbo P, Larson MG (2022) CutFEM based on extended finite element spaces. *Numer Math* 152:331–369
203. Badia S, Martín AF, Neiva E, Verdugo F (2021) The aggregated unfitted finite element method on parallel tree-based adaptive meshes. *SIAM J Sci Comput* 43(3):C203–C234
204. Hunt JD (2004) An adaptive three-dimensional Cartesian approach for the parallel computation of inviscid flow about static and dynamic configurations. PhD thesis, University of Michigan
205. Badia S, Neiva E, Verdugo F (2022) Robust high-order unfitted finite elements by interpolation-based discrete extension. *arXiv:2201.06632 [math.NA]*
206. Burman E, Hansbo P, Larson MG, Larsson K (2022) Extension operators for trimmed spline spaces. *arXiv:2205.06543 [math.NA]*
207. Badia S, Dilip H, Verdugo F (2022) Space–time aggregated finite element methods for time-dependent problems on moving domains. *arXiv:2206.03626 [math.NA]*
208. Verdugo F, Martín AF, Badia S (2019) Distributed-memory parallelization of the aggregated unfitted finite element method. *Computer Methods Appl Mech Eng* 357:112583, 2019. Figures reproduced with permission from Elsevier
209. Badia S, Martín AF, Principe J (2018) FEMPAR: an object-oriented parallel finite element framework. *Arch Comput Methods Eng* 25:195–271
210. Balay S, Abhyankar S, Adams MF, Brown J, Brune P, Buschelman K, Dalcin L, Dener A, Eijkhout V, Gropp WD, Kaushik D, Knepley MG, May DA, McInnes LC, Mills RT, Munson T, Rupp K, Sanan P, Smith BF, Zampini S, Zhang H, Zhang H (2019) PETSc users manual. Technical Report ANL-95/11 - Revision 3.12, Argonne National Laboratory
211. Badia S, Caicedo MA, Martín AF, Principe J (2021) A robust and scalable unfitted adaptive finite element framework for nonlinear solid mechanics. *Comput Methods Appl Mech Eng* 386:114093
212. Badia S, Martín AF, Verdugo F (2018) Mixed aggregated finite element methods for the unfitted discretization of the Stokes problem. *SIAM J Sci Comput* 40(6):B1541–B1576
213. Codina R, Badia S, Baiges J, Principe J (2017) Variational multiscale methods in computational fluid dynamics. *Encyclopedia of Computational Mechanics Second Edition*, pp 1–28
214. Dettmer WG, Kadapa C, Perić D (2016) A stabilised immersed boundary method on hierarchical b-spline grids. *Comput Methods Appl Mech Eng* 311:415–437
215. Hansbo P, Larson MG, Larsson K (2017) Cut finite element methods for linear elasticity problems. In Geometrically unfitted finite element methods and applications: Proceedings of the UCL Workshop 2016:578–593
216. Burman E, Hansbo P, Larson MG (2022) On the design of locking free ghost penalty stabilization and the relation to CutFEM with discrete extension. *arXiv:2205.01340 [math.NA]*
217. Preuss J (2018) Higher order unfitted isoparametric space–time FEM on moving domains. Master’s thesis, University of Göttingen
218. Lehrenfeld C, Olshanskii M (2019) An Eulerian finite element method for PDEs in time-dependent domains. *ESAIM: mathematical modelling and numerical analysis (ESAIM: M2AN)* 53(2):585–614
219. Larson MG, Zahedi S (2021) Conservative discontinuous cut finite element methods. *arXiv:2105.02202 [math.NA]*
220. Becker R, Braack M (2001) A finite element pressure gradient stabilization for the Stokes equations based on local projections. *Calcolo* 38:173–199

221. Badia S (2012) On stabilized finite element methods based on the Scott–Zhang projector. Circumventing the inf-sup condition for the Stokes problem. *Comput Methods Appl Mech Eng* 247–248:65–72
222. Johansson A, Larson MG, Logg A (2020) A multimesh finite element method for the Stokes problem. In *Numerical methods for flows: FEF 2017 selected contributions*, pp 189–198
223. Johansson A, Kehlet B, Larson MG, Logg A (2019) Multimesh finite element methods: solving PDEs on multiple intersecting meshes. *Comput Methods Appl Mech Eng* 343:672–689
224. Johansson A, Larson MG, Logg A (2020) Multimesh finite elements with flexible mesh sizes. *Comput Methods Appl Mech Eng* 372:113420
225. Nguyen L, Stoter SKF, Baum T, Kirschke J, Ruess M, Yosibash Z, Schillinger D (2017) Phase-field boundary conditions for the voxel finite cell method: Surface-free stress analysis of CT-based bone structures. *International Journal for Numerical Methods in Biomedical Engineering* 33(12):e2880
226. Hoang T, Verhoosel CV, Auricchio F, van Brummelen EH, Reali A (2018) Skeleton-stabilized isogeometric analysis: high-regularity interior-penalty methods for incompressible viscous flow problems. *Comput Methods Appl Mech Eng* 337:324–351
227. Guzmán J, Olshanskii M (2018) Inf-sup stability of geometrically unfitted Stokes finite elements. *Math Comput* 87:2091–2112
228. Hansbo P, Larson MG, Zahedi S (2015) Characteristic cut finite element methods for convection-diffusion problems on time dependent surfaces. *Comput Methods Appl Mech Eng* 293:431–461
229. Burman E, Hansbo P, Larson MG (2015) A stabilized cut finite element method for partial differential equations on surfaces: The Laplace–Beltrami operator. *Comput Methods Appl Mech Eng* 285:188–207
230. Clément P (1975) Approximation by finite element functions using local regularization. *Revue française d’automatique, informatique, recherche opérationnelle. Analyse numérique* 9(R2):77–84
231. Burman E, Hansbo P, Larson MG (2022) CutFEM based on extended finite element spaces. *arXiv:2101.10052 [math.NA]*
232. Brenner SC, Scott LR (2008) *The mathematical theory of finite element methods*. Springer, Berlin
233. Badia S, Martorell PA, Verdugo F (2022) Geometrical discretisations for unfitted finite elements on explicit boundary representations. *J Comput Phys* 460:111162
234. Fromm JE, Wunsch N, Xiang R, Zhao H, Maute K, Evans JA, Kamensky D (2023) Interpolation-based immersed finite element and isogeometric analysis. *Comput Methods Appl Mech Eng* 405:115890
235. Buffa A, Chanon O, Vázquez R (2022) Analysis-aware defeaturing: problem setting and a posteriori estimation. *Math Models Methods Appl Sci* 32(2):359–402
236. Zander N, Bog T, Kollmannsberger S, Schillinger D, Rank E (2015) Multi-level hp-adaptivity: high-order mesh adaptivity without the difficulties of constraining hanging nodes. *Comput Mech* 55:499–517
237. D’Angella D, Zander N, Kollmannsberger S, Frischmann F, Rank E, Schröder A, Reali A (2016) Multi-level *hp*-adaptivity and explicit error estimation. *Adv Model Simul Eng Sci* 3(33)
238. Burman E, He C, Larson MG (2020) A posteriori error estimates with boundary correction for a cut finite element method. *IMA J Numer Anal* 42(1):333–362
239. Divi SC, van Zuijlen PH, Hoang T, de Prenter F, Auricchio F, Reali A, van Brummelen EH, Verhoosel CV (2022) Residual-based error estimation and adaptivity for stabilized immersed isogeometric analysis using truncated hierarchical B-splines. *J Mech* 38:204–237
240. Becker R, Burman E, Hansbo P (2009) A Nitsche extended finite element method for incompressible elasticity with discontinuous modulus of elasticity. *Comput Methods Appl Mech Eng* 198(41):3352–3360
241. Gerstenberger A, Wall WA (2010) An embedded Dirichlet formulation for 3D continua. *Int J Numer Meth Eng* 82(5):537–563
242. Wolfe C (2009) The immersed boundary approach to fluid flow simulation, <https://www.ansys.com/resource-center/article/immersed-boundary-approach-to-fluid-flow-simulation-ansys-advantage-v3-i2>. Accessed 15 May 2022
243. Hartmann S, Leidinger LF, Benson DJ, Li L, Nagy AP, Pigazzini M (2019) Enabling the analysis of topologically connected multi-patch trimmed NURBS Shells in LS-DYNA. In: *Proceedings of the 12th European LS-DYNA conference*
244. Meßmer M, Leidinger LF, Hartmann S, Bauer F, Duddeck F, Wüchner R, Bletzinger K-U (2021) Isogeometric analysis on trimmed solids: a B-spline-based approach focusing on explicit dynamics. In: *Proceedings of the 13th European LS-DYNA conference*
245. Favaloro AJ, Brenken B, Barocio E, Pipes RB (2017) Simulation of polymeric composites additive manufacturing using Abaqus. In: *Proceedings of Dassault Systèmes’ science in the age of experience*, pp 103–114
246. Courter B, Savane V, Bi J, Dev S, Hansen CJ (2017) Finite element simulation of the fused deposition modelling process. In: *Proceedings of the NAFEMS world Ccongress*, pp 11–14
247. Hyperganic. <https://www.hyperganic.com>. Assessed 15 May 2022
248. Bangerth W, Hartmann R, Kanschat G (2007) deal.II—a general-purpose object-oriented finite element library. *ACM Trans Math Softw* 33(4):Article 24
249. Alnæs M, Blechta J, Hake J, Johansson A, Kehlet B, Logg A, Richardson C, Ring J, Rognes ME, Wells GN (2015) The FEniCS project version 1.5. *Arch Numer Softw* 3(100):9–23
250. Badia S, Verdugo F (2020) Gridap: An extensible Finite Element toolbox in Julia. *J Open Source Softw* 5(52):2520
251. Schöberl J (1997) NETGEN An advancing front 2D/3D-mesh generator based on abstract rules. *Comput Vis Sci* 1:41–52
252. Bastian P, Blatt M, Dedner A, Dreier N-A, Engwer C, Fritze R, Gräser C, Grüninger C, Kempf D, Klöforn R, Ohlberger M, Sander O (2021) The DUNE framework: basic concepts and recent developments. *Comput Math Appl* 81:75–112

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.