

**Enriched immersed finite element and isogeometric analysis
algorithms and data structures**

Wunsch, Nils; Doble, Keenan; Schmidt, Mathias R.; Noël, Lise; Evans, John A.; Maute, Kurt

DOI

[10.1007/s00366-025-02163-7](https://doi.org/10.1007/s00366-025-02163-7)

Publication date

2025

Document Version

Final published version

Published in

Engineering with Computers

Citation (APA)

Wunsch, N., Doble, K., Schmidt, M. R., Noël, L., Evans, J. A., & Maute, K. (2025). Enriched immersed finite element and isogeometric analysis: algorithms and data structures. *Engineering with Computers*, 41(6), 3919-3957. <https://doi.org/10.1007/s00366-025-02163-7>

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.



Enriched immersed finite element and isogeometric analysis: algorithms and data structures

Nils Wunsch¹ · Keenan Doble¹ · Mathias R. Schmidt² · Lise Noël³ · John A. Evans¹ · Kurt Maute¹

Received: 18 December 2024 / Accepted: 25 May 2025 / Published online: 18 August 2025
© The Author(s) 2025

Abstract

Immersed finite element methods provide a convenient analysis framework for problems involving geometrically complex domains, such as those found in topology optimization and microstructures for engineered materials. However, their implementation remains a major challenge due to, among other things, the need to apply nontrivial stabilization schemes and generate custom quadrature rules. This article introduces the robust and computationally efficient algorithms and data structures comprising an immersed finite element preprocessing framework. The input to the preprocessor consists of a background mesh and one or more geometries defined on its domain. The output is structured into groups of elements with custom quadrature rules formatted such that common finite element assembly routines may be used without or with only minimal modifications. The key to the preprocessing framework is the construction of material topology information, concurrently with the generation of a quadrature rule, which is then used to perform enrichment and generate stabilization rules. While the algorithmic framework applies to a wide range of immersed finite element methods using different types of meshes, integration, and stabilization schemes, the preprocessor is presented within the context of the extended isogeometric analysis. This method utilizes a structured B-spline mesh, a generalized Heaviside enrichment strategy considering the material layout within individual basis functions' supports, and face-oriented ghost stabilization. Using a set of examples, the effectiveness of the enrichment and stabilization strategies is demonstrated alongside the preprocessor's robustness in geometric edge cases. Additionally, the performance and parallel scalability of the implementation are evaluated.

Keywords Immersed finite element method · XIGA · Heaviside enrichment · Multi-material problems · Ghost stabilization · Computer implementation

Nils Wunsch, Keenan Doble contributed equally to this work.

✉ Kurt Maute
kurt.maute@colorado.edu

Nils Wunsch
nils.wunsch@colorado.edu

Keenan Doble
keenan.doble@colorado.edu

Mathias R. Schmidt
schmidt43@llnl.gov

Lise Noël
l.f.p.noel@tudelft.nl

John A. Evans
john.a.evans@colorado.edu

¹ Smead Aerospace Engineering Sciences, University of Colorado Boulder, 3775 Discovery Dr., Boulder 80303, CO, USA

List of symbols

Numerical symbols: operators & conventions

$(\cdot)^h$	Discrete quantity
(\cdot)	Enriched function (space)
$\overline{(\cdot)}$	Polynomial extension
$\overline{(\cdot)}$	Closure of a domain
$\partial(\cdot)$	Boundary of a domain
$\{\cdot\}$	Weighted interface operator, see (8)
$[[\cdot]]$	Jump operator, see (9)
$\ \cdot\ _{L2}$	L2 norm
$ \cdot _{H1}$	H1 semi-norm

² Computational Engineering Division, Lawrence Livermore National Laboratory, 7000 East Ave., Livermore 94550, CA, USA

³ Department of Precision and Microsystems Engineering, Delft University of Technology, Mekelweg 2, Delft 2628, CD, The Netherlands

$\partial_{\mathbf{n}}^{(i)}$	i -th derivative in normal direction \mathbf{n}	CM	Child mesh
		d	Dimensionality
		D	Cluster
Numerical quantities		e	Entity (index)
\mathbf{b}	Volumetric body load	ε	Enrichment level
$\mathcal{B}, \mathcal{B}_E$	Finite element basis, local to the bg. element	E	Background element (index)
		F	Facet (index)
d, \mathbf{d}	Degree of freedom, solution vector	G	Geometry
E	Young's modulus	g	Geometry index
h, h_0	(Initial) background element size	\mathcal{H}	Background mesh
\mathbf{n}, \mathbf{n}_m	Normal, pointing out from Ω_m	i, j, k, l	General indices
N	Basis function	I	Parallel ID
p	Polynomial degree	ℓ	Enriched B.F. multi-index (B, ε)
$\mathcal{R}_{(\cdot)}$	Residual for (B-bulk, D-Dirichlet, N-Neumann, I-interface, G-ghost)	m	Material (index)
		$n_{(\cdot)}$	Number of (\cdot)
\mathbf{t}	Traction vector	o	Ordinal
\mathbf{u}	State variable/displacement field	P	Proximity
\mathbf{u}_0	Prescribed displacement at boundary	p	(Current) processor ID
\mathbf{v}	Test function/virtual displacement	q	Quadrature point index
\mathcal{V}	Space of test/trial functions	Q_v, Q_c	Vertex cell queues
w	Interface weight, see (8)	r	(Entity) rank
\mathbf{x}	Physical coordinates	S	Subphase
$\gamma_{(\cdot)}$	Penalty parameters (D-Dirichlet, I-interface, G-ghost)	\mathcal{T}	Foreground mesh
		u, v	Unzipping index
$\boldsymbol{\epsilon}$	Strain vector/tensor	v	Vertex (index)
κ	Condition number		
$\boldsymbol{\sigma}$	Stress tensor		
ν	Poisson's ratio		
ψ	Enrichment function, see (12)		
ξ	Parametric coordinates		
Domains		Sets and graphs	
\mathcal{A}	Ambient domain (domain of the background mesh)	\mathcal{C}	Set of all fg. cells
\mathcal{F}_G	Skeleton domain for face-oriented ghost stabilization, see (18)	$\mathcal{D}_{(\cdot)}$	Set of all clusters of type (B-bulk, S-side, I-interface, G-ghost)
R, R_B^{ε}	Connected domain within B.F. support, support of the enriched B.F. $\tilde{N}_B^{\varepsilon}$, see Fig. 3	\mathcal{E}	Set of all bg. elements
S	Connected domain within bg. element	\mathcal{F}	Facet connectivity (graph)
Γ^D, Γ^N	Dirichlet Neumann boundaries	$\mathcal{G}_S, \mathcal{G}_I$	Graph of connected, disconnected subphases
$\Gamma_{m,l}^I$	Material interface between materials with indices m and l	\mathcal{M}	Set of (non-void) material indices
$\overset{\circ}{\Gamma}^I$	Union of material interfaces in the interior of \mathcal{A} , see (17)	\mathcal{S}	Set of all subphases
Ω_m	Material domain	$\mathcal{U}_{u,E}$	Set of subphases connected to S_E^u , see (20)
Ω^E	Domain of a background element	CtE	Cell-to-entity connectivity, see Fig. 9
$\overset{\circ}{\Omega}$	Union of material domains, see (1)	DF	Sets of descendant facets, see Algorithm 4
Implementation: indices and entities/objects		EtC	Entity-to-cell connectivity, see Fig. 9
a	Ancestor entity (index)	IEN	Local to global basis function map
b, B	Local global basis function indices	NVI	Set of new vertex indices
c	Cell (index)	P_{glob}	Global communication table (set of all processors)
		P_{loc}	Local communication table (set of neighboring processors)
		Ξ	Set of unzipped bg. elements

Compute performance

M	Memory consumption
n_p	Number of processors
s	Problem size
t	Compute time
η	Compute efficiency, see (25)
μ	Memory efficiency, see (26)
λ_{loc}	Local mesh overhead, see (27)
λ_{glob}	Global mesh overhead, see (28)
$\mathcal{O}(\cdot)$	Order of limiting behavior
$(\cdot)_0$	Value in reference run
$(\cdot)'$	Relative quantity = $(\cdot)/(\cdot)_0$

Contents

1	Introduction	3921
2	Numerical framework	3923
2.1	Discretization	3923
2.2	Element formation and integration	3925
2.3	Face-oriented ghost stabilization	3926
3	Mesh generation	3927
3.1	Preliminaries	3927
3.2	Conformal foreground mesh generation	3930
3.2.1	Regular subdivision	3932
3.2.2	Templated subdivision	3933
3.3	Generation of topological information	3935
3.4	Enrichment	3937
3.5	Face-oriented ghost stabilization	3940
3.6	Parallelization	3941
4	Numerical examples	3943
4.1	Multiple thin beams	3944
4.2	Brick wall	3945
4.3	Scaffold sandwich	3947
5	Conclusion	3950
Appendix A	Glossary	3951
Appendix B	Additional Algorithms and Data Structures	3952
B.1	Determining Ancestry of New Vertices	3952
B.2	Entity Connectivity	3952
B.3	Foreground Mesh – Complete Data Structure	3954
References		3954

1 Introduction

The finite element method (FEM) [42] has been adopted as a standard tool to solve continuum mechanics problems in engineering analysis. In its traditional form, the FEM requires boundary-fitted meshes consisting of elements that define both the approximation of the state variable fields, e.g. temperature or displacement, and the geometric domain over which the weak statement of the governing partial differential equations (PDEs) is integrated. To obtain high-fidelity solutions, the finite element (FE) function spaces need to be of sufficient quality which may result in nontrivial geometric constraints on the elements comprising the mesh. For complex domain shapes, such as the fiber composite shown in



Fig. 1 Fiber composite weave with two types of fibers and a binding matrix material. The matrix material is shown in off-white and partially removed for better visibility

Fig. 1, boundary-fitted meshes satisfying these quality measures can generally not be generated reliably without manual intervention. This renders mesh generation a major bottleneck in computational engineering analysis [10, 21].

"Unfitted" [36] or "immersed" [94] FEMs have gained significant traction in recent years as an elegant solution to automate finite element analysis. The commonly used methods for immersed finite element analysis include the CutFEM [16] (which originated from partition of unity (PU) based methods [19], namely, the PU-FEM [5, 54]), the generalized FEM [22, 23, 82, 83], and the extended FEM [11, 55]. Other lineages of immersed methods include the finite cell method [24, 63, 73] (which originated from fictitious domain methods) and the shifted boundary method [50, 51]. These methods are specifically designed to accommodate arbitrary strong or weak discontinuities within elements, representing features such as cracks, domain boundaries, or material interfaces. As a result, a mesh, often referred to as the "background mesh," can be constructed on a geometrically simple domain into which a geometrically complex domain is immersed. The immersed approach is of particular interest and has been successfully applied to problems where the meshing effort is exacerbated by complex or evolving interface geometries such as in scan-based analysis [26, 67, 89], topology optimization [33, 60, 61, 88, 90], or physical phenomena such as multi-phase flows [69] and fluid–structure interaction [40, 41, 46].

While immersed methods provide a convenient and time-saving design-through-analysis workflow, the approach departs from traditional FEMs in four key aspects. To illustrate these non-standard aspects, consider the cross-section of a fiber composite immersed into a background mesh, as shown in Fig. 2.

- (i) Multiple materials, components, or other features like cracks may be present within a cut element, as shown

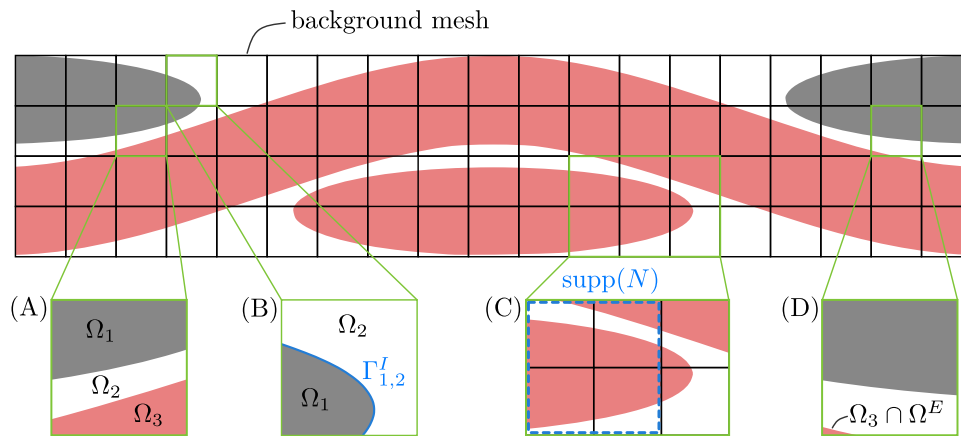


Fig. 2 Cross-section of the fiber weave shown in Fig. 1 with a coarse background mesh. Details: **A** Cut background element occupied by three different materials. **B** Non-trivially intersected background element with highlighted material interface $\Gamma_{1,2}^I$. **C** Sub-element scale

material features in neighboring elements. Highlighted is a single basis function's support $\text{supp}(N)$ which spans multiple features. **(D)** Cut background element where one material subdomain occupies only a small volume fraction

in Fig. 2 detail (A). The finite element basis needs to resolve the associated discontinuities inside the background elements. This is commonly done through basis function enrichment [28, 79].

- (ii) Essential boundary and interface conditions cannot be built into the test and trial function spaces. Instead, these conditions require weak enforcement through the penalty method [96], Lagrange multiplier methods [13], or Nitsche's method [14]. For an overview of the issue, the reader is referred to [27].
- (iii) The governing equations must be satisfied on given material domains Ω_m . However, the background elements do not coincide with these material regions. As a result, these elements cannot be reused to generate quadrature rules to evaluate the integrals of the associated weak form. Instead, custom quadrature rules need to be constructed for each background element intersected by a material interface $\Gamma_{m,l}^I$ and the material subdomains within them. Furthermore, quadrature rules must be constructed for segments of the interface $\Gamma_{m,l}^I$ within background elements, such as the one in Fig. 2 detail (B), to evaluate the integrals enforcing the boundary conditions.
- (iv) Lastly, the arbitrary shape and location of an interface within a background element may lead to small volume fractions $\text{vol}(\Omega_m \cap \Omega^E) \ll \text{vol}(\Omega^E)$ within that background element's domain Ω^E , as shown in Fig. 2 detail (D). Such configurations can lead to stability and conditioning issues and require stabilization. For an overview of this issue, the reader is referred to [65].

The generation of quadrature rules and the construction of the enriched finite element space, as well as the application of appropriate stabilization schemes, require

extensive supplemental geometric information that needs to be generated during preprocessing. These requirements lead to a challenging implementation of immersed methods. Multiple approaches have been developed to address each of the four aspects individually. However, the literature mostly presents these methods from a theoretical standpoint or covers the implementation of individual methods independently of how other aspects are addressed. The implementation of immersed finite elements and the associated methods is indirectly addressed by publicly available source code. Examples of published open-source FE packages with immersed FE capabilities are `ngsxfem` [48], `deal.II` [3], `MOOSE-XFEM` [44, 95], and the discontinued libraries `LibCutFEM` [16] and `MultiMesh` [45]. The works by Zhang et al. [93] and Badia et al. [6] are an exception to this. The former provides an overview of the algorithms and data structures underlying an enriched/immersed FE preprocessor and discusses the interaction of various components of the framework. However, it does not cover the stabilization aspect or the realization of a parallel implementation. Similarly, the latter presents a detailed implementation of the package `FEMPAR` capable of immersed FE discretizations, including a discussion of stabilization and parallel processing; but the implementation is still tied to a specific code base, limiting its applicability in other contexts, such as multi-material problems requiring basis enrichment, which is not supported by it.

The authors believe that there is significant value in the discussion of the implementation of an immersed FE framework, as not only is the translation from theory into code challenging, but so is doing so efficiently.

This article presents a comprehensive implementation of a preprocessor that generates all data necessary to perform immersed finite element analysis on multi-material prob-

lems, including the construction of enriched finite element spaces, as well as the generation of custom quadrature and stabilization rules. We discuss the algorithmic design choices made and their impact on robustness and performance characteristics, specifically runtime, memory consumption, and parallel scalability. A further focus is placed on how the four key aspects described previously are addressed cohesively. Although the implementation is presented in the context of the extended isogeometric analysis (XIGA) introduced by [62] and [75], the algorithms and data structures are general and adaptable to suit other immersed analysis frameworks that may utilize different types of background meshes, geometry descriptions, or use alternative methods to generate quadrature and stabilization rules. The preprocessor can also be used for interpolation-based immersed analysis [30, 31]. The implementation, as presented in this paper, is part of the immersed FE and optimization package *Moris*, whose C++ source code is also [publicly available](#) [1]. The framework has been developed with a focus on solving multi-material level-set (LS) topology optimization problems, where the shape of objects and their material composition are defined implicitly by a set of LS functions. Such functions can be created, e.g., from images or 3D scans or via geometry modeling tools using function representations as done by *nTop* [91] and *openVCAD* [92].

The XIGA employs a generalized Heaviside enrichment strategy in combination with Nitsche's method [59] to resolve material interfaces and the associated discontinuities inside background elements. Most Heaviside enrichment strategies consider either the material domains or their globally connected subdomains [37, 87], i.e., material components, to construct the enriched function space. However, configurations with features on a length scale smaller than a background element may arise, e.g., in the context of topology optimization or engineered materials. This leads to an artificial numerical coupling, sometimes referred to as “cross-talk” [49], due to basis functions that span features of the same globally connected material domain but are disconnected within that function's support, as shown in Fig. 2 detail (C). The issue is exacerbated when basis functions with larger supports are employed, such as high-continuity splines. The literature on isogeometric analysis has addressed this issue through either sufficient local refinement [20] or modifications to the basis via control-point insertion [49] near non-convex features. To mitigate cross-talk in this framework, the generalized Heaviside enrichment strategy considers the material layout on the level of the individual basis function supports and appropriately creates additional degrees of freedom, where artificial coupling occurs [62].

The information about the material layout necessary to apply the generalized Heaviside enrichment strategy is obtained by building connectivity graphs after the tessellation of the background mesh. The tessellation procedure

employed relies on the application of subdivision templates as done by, e.g., [81]. Custom quadrature rules on the various material subdomains are generated by mapping a Gauss-Legendre quadrature rule from the resulting boundary-fitted “foreground” mesh onto the background elements. Alternative, potentially higher-order accurate, quadrature methods and tessellation procedures may be employed, such as, e.g., the ones presented in [8] and [52] for explicit geometry descriptions or the one in [29] for implicit geometry descriptions, so long as the necessary arrangements to build a material connectivity graph are made.

Lastly, the stabilization aspect can be addressed using methods such as basis function removal [25], basis function extension approaches [7, 17, 18], or ghost stabilization [12, 15]. With the information about the material layout within intersected background elements previously obtained, either of these stabilization methods can be applied with relative ease. For this work, a face-oriented ghost stabilization [78] is chosen, as, unlike the other methods, the ghost stabilization does not rely on non-intersected elements in the vicinity of an interface, resulting in better robustness in pathological geometric configurations.

The remainder of this article is structured as follows. First, Sect. 2 lays out the general challenges of an enriched immersed method and the mathematical formulations of the enrichment and ghost stabilization strategies. Using this knowledge, the output information, which the preprocessor needs to generate, is formalized. Section 3 covers the algorithms and data structures underlying the implementation of the preprocessing framework. Lastly, Sect. 4 demonstrates robustness, efficiency, and parallel scalability using a suite of examples.

2 Numerical framework

This section discusses the theoretical aspects of the computational framework and focuses on each of the framework's components, which were first introduced in [62], with the intent of providing the requirements for its implementation. Some details are omitted for brevity; the reader is referred to [62] for a detailed exposure to the underlying theory. We start the discussion by first presenting the discretized weak form for the fiber patch shown in Fig. 1 as a linear elastic model problem and then address each of the aspects necessary for constructing the information to assemble the system of equations within the given framework.

2.1 Discretization

Model Problem and Weak Form

The fiber patch, shown in Fig. 1, is composed of three materials $m \in \mathcal{M} = \{1, 2, 3\}$ (red fibers, gray fibers, and

matrix material), all modeled as linear elastic, but with different material parameters. The displacements at the interfaces are assumed to be matching. The patch is subjected to external loads and displacement constraints. For convenience, we define the union $\hat{\Omega}$ of domains occupied by each material Ω_m , i.e., the physical domain without the internal interfaces, as

$$\hat{\Omega} = \bigcup_{m \in \mathcal{M}} \Omega_m. \tag{1}$$

The discretized weak form for the problem with discrete trial and test displacement \mathbf{u}^h and \mathbf{v}^h states as: find $\mathbf{u}^h \in \mathcal{V}^h \subset H^1(\hat{\Omega})$ such that, $\forall \mathbf{v}^h \in \mathcal{V}^h \subset H^1(\hat{\Omega})$:

$$\begin{aligned} \mathcal{R}_\Omega(\mathbf{u}^h, \mathbf{v}^h) + \mathcal{R}_N(\mathbf{u}^h, \mathbf{v}^h) + \mathcal{R}_D(\mathbf{u}^h, \mathbf{v}^h) + \dots \\ \dots + \mathcal{R}_I(\mathbf{u}^h, \mathbf{v}^h) + \mathcal{R}_G(\mathbf{u}^h, \mathbf{v}^h) = 0, \end{aligned} \tag{2}$$

where

$$\mathcal{R}_\Omega(\mathbf{u}^h, \mathbf{v}^h) = \sum_{m \in \mathcal{M}} \int_{\Omega_m} \boldsymbol{\epsilon}(\mathbf{v}^h) : \boldsymbol{\sigma}(\mathbf{u}^h) - \mathbf{b} \, d\Omega, \tag{3}$$

$$\mathcal{R}_N(\mathbf{u}^h, \mathbf{v}^h) = \sum_{m \in \mathcal{M}} \int_{\Gamma_m^N} -\mathbf{v}^h \cdot \mathbf{t} \, d\Gamma, \tag{4}$$

$$\begin{aligned} \mathcal{R}_D(\mathbf{u}^h, \mathbf{v}^h) \\ = \sum_{m \in \mathcal{M}} \left[- \int_{\Gamma_m^D} \mathbf{v}^h \cdot (\boldsymbol{\sigma}(\mathbf{u}^h) \cdot \mathbf{n}_m) \, d\Gamma \right. \\ \dots \mp \int_{\Gamma_m^D} (\boldsymbol{\sigma}(\mathbf{v}^h) \cdot \mathbf{n}_m) \cdot (\mathbf{u}^h - \mathbf{u}_0) \, d\Gamma \\ \left. \dots + \int_{\Gamma^D} \gamma_D \mathbf{v}^h \cdot (\mathbf{u}^h - \mathbf{u}_0) \, d\Gamma \right], \end{aligned} \tag{5}$$

$$\begin{aligned} \mathcal{R}_I(\mathbf{u}^h, \mathbf{v}^h) \\ = \sum_{\substack{m, l \in \mathcal{M} \\ m \neq l}} \left[- \int_{\Gamma_{m, l}^I} \llbracket \mathbf{v}^h \rrbracket \cdot \{ \boldsymbol{\sigma}(\mathbf{u}^h) \cdot \mathbf{n}_m \} \, d\Gamma \right. \\ \dots \mp \int_{\Gamma_{m, l}^I} \{ \boldsymbol{\sigma}(\mathbf{v}^h) \cdot \mathbf{n}_l \} \cdot \llbracket \mathbf{u}^h \rrbracket \, d\Gamma \\ \left. \dots + \int_{\Gamma_{m, l}^I} \gamma_I \llbracket \mathbf{v}^h \rrbracket \cdot \llbracket \mathbf{u}^h \rrbracket \, d\Gamma \right]. \end{aligned} \tag{6}$$

(3) and (4) are the bulk and Neumann terms, originating from the strong form and integration by parts. Here $\boldsymbol{\epsilon}$, $\boldsymbol{\sigma}$, \mathbf{t} , and \mathbf{b} denote the strain, stress, prescribed traction, and volumetric body force, respectively; \mathbf{n}_m denotes the outward unit normal to the interface associated with the material domain Ω_m . The terms used to weakly enforce the Dirichlet boundary conditions via Nitsche’s method [14, 59] are collected in (5). The contribution (6) weakly enforces matching displacements and tractions at interfaces between materials

ments and tractions at interfaces between materials

$$\Gamma_{m, l}^I = \partial\Omega_m \cap \partial\Omega_l \text{ for } m, l \in \mathcal{M}, m \neq l, \tag{7}$$

where $\partial(\cdot)$ denotes the boundary of a domain. The braces and double brackets are jump terms defined as

$$\{f(\mathbf{x})\} = \lim_{\delta \rightarrow 0^+} w_m f(\mathbf{x} - \delta \mathbf{n}_m) + w_l f(\mathbf{x} + \delta \mathbf{n}_m), \tag{8}$$

$$\llbracket f(\mathbf{x}) \rrbracket = \lim_{\delta \rightarrow 0^+} f(\mathbf{x} - \delta \mathbf{n}_m) - f(\mathbf{x} + \delta \mathbf{n}_m), \tag{9}$$

with $w_m + w_l = 1$ and $\mathbf{x} \in \Gamma_{m, l}^I$. The choice of weights w_m and w_l , as well as the penalty parameters γ_D and γ_I are adopted from [2]. Further, the signs of the symmetry terms in (5) and (6) are chosen to yield either a symmetric or a non-symmetric formulation. We briefly discuss the consequences of this choice in Sect. 2.3 alongside the ghost term \mathcal{R}_G .

Note that every material subdomain Ω_m is treated separately to allow for, e.g., the use of different constitutive laws or physics for each material, or allow for a wider range of interface conditions to be enforced via the term \mathcal{R}_I . Materials may be chosen to be void, thus not contributing to the residual, and some interfaces $\Gamma_{m, l}^I$ become part of the Dirichlet or Neumann boundaries.

Approximation Spaces and Enrichment

The finite element space $\mathcal{V}^h = \text{span}(\mathcal{B})$ is provided by a set of basis functions $\mathcal{B} = \{N_B(\mathbf{x})\}_{B=1}^{n_B}$ with

$$\mathbf{u}^h(\mathbf{x}) = \sum_{B=1}^{n_B} N_B(\mathbf{x}) \, d_B, \tag{10}$$

where d_B are the weights associated with each basis function N_B . Throughout the remainder of this paper, we are using B to index basis functions. Further, $n_{(\cdot)}$ denotes the number of whatever the subscript index is associated with. The basis functions are defined on the elements of the background mesh $\mathcal{H} = \{\Omega^E\}_{E=1}^{n_E}$ which does not conform to interfaces, and we assume the basis functions to be compactly supported by a set of background elements, without making any further assumptions on the type of basis function used at this time.

To restrict the support of basis functions to the material domains and allow for the enforcement of arbitrary interface conditions, a Heaviside enrichment strategy is employed. The specific strategy used in this framework considers different material subdomains within a single basis function’s support, as illustrated in Fig. 3. This prevents artificial numerical coupling between small geometry features without the need for local refinement, as demonstrated in Sect. 4.1.

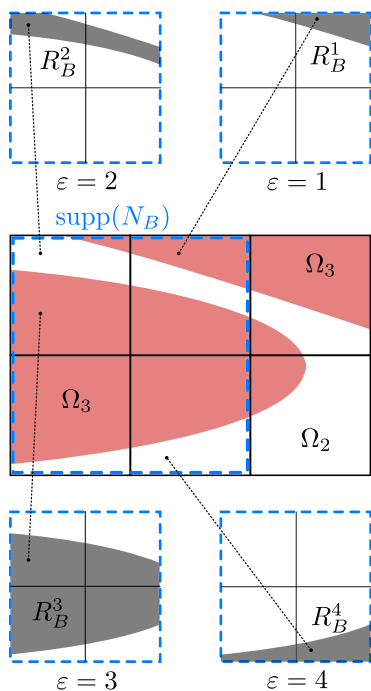


Fig. 3 Enrichment strategy for a single basis function for the material layout shown in Fig. 2, detail (C). The basis function’s support is shown in dashed blue lines. The support contains four connected subdomains R_B^ϵ for each of which an enriched basis function is defined, where $\tilde{N}_B^\epsilon = \psi_B^\epsilon N_B$ with $\text{supp}(\tilde{N}_B^\epsilon) = R_B^\epsilon$

With this enrichment strategy, the approximation of the state variables reads as follows:

$$\mathbf{u}^h(\mathbf{x}) = \sum_{B=1}^{n_B} \sum_{\epsilon=1}^{n_\epsilon(B)} \psi_B^\epsilon(\mathbf{x}) N_B(\mathbf{x}) d_B^\epsilon, \tag{11}$$

where we define the indicator, or "enrichment", function ψ_B^ϵ as

$$\psi_B^\epsilon(\mathbf{x}) = \begin{cases} 1, & \mathbf{x} \in R_B^\epsilon, \\ 0, & \text{otherwise.} \end{cases} \tag{12}$$

Here $R_B^\epsilon \subseteq \text{supp}(N_B)$ with $\bigcup_{\epsilon=1}^{n_\epsilon(B)} R_B^\epsilon = \text{supp}(N_B) \cap \hat{\Omega}$ correspond to the disconnected subdomains within the basis function’s support, as is shown in Fig. 3.

We refer to the product $\tilde{N}_B^\epsilon = \psi_B^\epsilon \cdot N_B$ as an "enriched basis function". The set of all enriched basis functions is defined as

$$\tilde{\mathcal{B}} := \{ \psi_B^\epsilon(\mathbf{x}) N_B(\mathbf{x}) : B \in \{1, \dots, n_B\}, \epsilon \in \{1, \dots, n_\epsilon(B)\} \} \tag{13}$$

From this arises the requirement on the implementation to identify the number of enrichment levels $n_\epsilon(B)$ for each basis function N_B and construct their enriched counterparts \tilde{N}_B^ϵ . As is detailed in Sect. 3.4, a flood-fill algorithm is used to identify

the disconnected subdomains. This approach subsequently requires that the material connectivity is available.

2.2 Element formation and integration

In a body-fitted finite element approach, each finite element has a set of non-zero basis functions supported within it and belongs to part of a material domain. The integrals (3) - (6) can hence be written as a sum of integrals over each element. The evaluation of these elemental integrals is referred to as "element formation". In contrast, for an immersed approach with multiple material domains, the elemental integrals need to be further split into the different material or interface subdomains within a given background element. Taking into account the enrichment strategy (11), we specifically consider materially disconnected subdomains within a background element – rather than material subdomains – such that it can be ensured that each subdomain supports a set of non-zero enriched basis functions. The generalized form of the bulk, boundary, and interface integrals (3) - (6) can then be stated as

$$\sum_{m \in \mathcal{M}} \int_{\Omega_m} (\cdot) d\Omega = \sum_{E=1}^{n_E} \sum_{u=1}^{n_u(E)} \int_{S_E^u} (\cdot) d\Omega, \tag{14}$$

$$\sum_{m \in \mathcal{M}} \int_{\Gamma_m} (\cdot) d\Omega = \sum_{E=1}^{n_E} \sum_{u=1}^{n_u(E)} \int_{\partial S_E^u \cap \Gamma} (\cdot) d\Gamma, \tag{15}$$

$$\sum_{\substack{m,l \in \mathcal{M} \\ m \neq l}} \int_{\Gamma_{m,l}} (\cdot) d\Omega = \sum_{E=1}^{n_E} \sum_{\substack{u,v=1 \\ u \neq v}}^{n_u(E)} \int_{\partial S_E^u \cap \partial S_E^v} (\cdot) d\Gamma. \tag{16}$$

The background elements are indexed by E , and we correspondingly denote the number of background elements by n_E . The number of disconnected subdomains S_E^u within a given background element’s domain Ω^E is denoted by $n_u(E)$. From an implementation standpoint, the challenge of evaluating the integrals is twofold. The first challenge is to identify the disconnected subdomains $\{S_E^u\}_{u=1}^{n_u(E)}$ within each background element E and the enriched basis functions supported within each subphase S_E^u . The second challenge is to construct custom quadrature rules for each subdomain to evaluate the integrals in (14) - (16).

Element formation in the immersed setting is generalized for any type of integral by introducing the concept of a cluster, as shown in Fig. 4. We define a cluster to be a background element with domain Ω^E , a single subdomain S_E^u , the set of enriched basis functions supported inside that subdomain, and a quadrature rule associated with either that subdomain’s volume or any part of its boundary ∂S_E^u . The quadrature rules themselves consist of a set of quadrature points and associated weights. In the case of interface integrals, a linked pair of clusters is needed to evaluate the jump terms (8) and (9).




Fig. 4 Concept of a cluster: a background element with a single active subdomain S_E^u , a set of enriched basis functions supported inside that subdomain, and a quadrature rule associated with either that subdomain (A, B) or any part of the subdomain's boundary ∂S_E^u (C). Blue indicates

the volume or surface integrated over; red crosses represent exemplary quadrature points. Additionally, linked pairs of clusters are constructed for interfaces (D)

The step of element formation is then reduced to evaluating integrands using the basis defined on the cluster (pair)'s background element(s) at the quadrature points and summing up the contributions.

Within the presented framework, the disconnected subdomains are obtained by a tessellation of intersected background elements using a fast, templated subdivision scheme detailed in Sect. 3.2.

2.3 Face-oriented ghost stabilization

Small material slivers lead to ill-conditioning due to the vanishing contributions of sparsely supported basis functions to the global system of equations [65]. Further, stability issues may arise due to the penalty terms in (5) and (6) not being any longer bounded for arbitrarily small material fractions.¹

The framework employs a version of the face-oriented ghost stabilization proposed by [15] adapted to the enrichment strategy used. The basic idea underlying face-oriented ghost stabilization is to penalize the jumps in derivatives across elemental faces adjacent to intersected background elements, thereby ensuring a sufficient contribution of sparsely supported basis functions. The penalties are only applied to those faces across which connected material regions transition, to not introduce spurious numerical coupling.

To illustrate the adaptations necessary for the enrichment strategy, consider the material layout shown in Fig. 5 (C). The background elements E^+ and E^- contain multiple material subdomains $\{S_{E^+}^u\}_{u=1}^{n_u(E^+)}$ and $\{S_{E^-}^u\}_{u=1}^{n_u(E^-)}$ with different sets of enriched basis functions supported within each sub-

domain. Additionally, the subdomains $S_{E^+}^1$ and $S_{E^+}^3$ support different sets of enriched basis functions but are part of the same material domain Ω_1 connected in the neighboring element E^- . Hence, the penalty needs to be applied to two different pairs of element bases on the same facet to involve all sparsely supported basis functions and appropriately stabilize the problem.

The procedure is generalized and formalized as follows. Let's define the union $\hat{\Gamma}^I$ of all material interfaces and boundaries inside the ambient domain \mathcal{A} as

$$\hat{\Gamma}^I = \bigcup_{\substack{m,l \in \mathcal{M} \\ m \neq l}} \Gamma_{m,l}^I. \quad (17)$$

The set of all ghost faces is defined as the set of faces F between neighboring background elements Ω^{E^+} and Ω^{E^-} , at least one of which is intersected by an interior interface $\hat{\Gamma}^I$

$$\mathcal{F}_G = \left\{ F : \Omega^{E^+} \cap \hat{\Gamma}^I \neq \emptyset \vee \Omega^{E^-} \cap \hat{\Gamma}^I \neq \emptyset \right. \\ \left. \overline{F} = \overline{\Omega^{E^+}} \cap \overline{\Omega^{E^-}}, E^+ \neq E^- \right\}, \quad (18)$$

where $\overline{(\cdot)}$ denotes the completion of a domain.

The contribution to the weak form (2) then reads as:

$$\mathcal{R}_G(\mathbf{u}^h, \mathbf{v}^h) = \sum_{F \in \mathcal{F}_G} \sum_{u=1}^{n_u(E^+)} \sum_{v \in \mathcal{U}_{u,E^-}} \dots \\ \dots \sum_{k=1}^P \int_F \gamma_G^k \llbracket \partial_{\mathbf{n}}^k \tilde{\mathbf{v}}^h \rrbracket \cdot \llbracket \partial_{\mathbf{n}}^k \tilde{\mathbf{u}}^h \rrbracket d\Gamma \quad (19)$$

where \mathcal{U}_{u,E^-} is the set of all subdomains in Ω^{E^-} that are connected to $S_{E^+}^u$ across the facet F defined as follows:

$$\mathcal{U}_{u,E^-} = \left\{ v \in \{1, \dots, n_u(E^-)\} : \dots \right. \\ \left. \dots m(S_{E^+}^u) = m(S_{E^-}^v) \wedge \left| \overline{S_{E^+}^u} \cap \overline{S_{E^-}^v} \right| > 0 \right\}, \quad (20)$$

¹ The stability parameters γ in (5) and (6) need to scale inversely with elemental material volumes (which can be arbitrarily small) to retain coercivity and, hence, well-posedness for the symmetric Nitsche formulation [64]. An asymmetric version of Nitsche's method avoids this issue at the cost of losing symmetry in the linear system of equations and adjoint consistency of the bilinear form resulting in potentially reduced L^2 -convergence [74].




Fig. 5 Ghost faces for an example material layout. **A** Material layout of a two-material problem with materials $\mathcal{M} = \{1, 2\}$ occupying a domain smaller than the ambient domain \mathcal{A} filled by the background mesh. The

white domain Ω_0 is assumed to be void. **B** The set of interior element faces \mathcal{F}_G which the ghost penalty is applied to. **C** Exemplary ghost facet F and the material layout in its adjacent background elements




Fig. 6 Ghost cluster pairs constructed for the example ghost facet and material layout shown in Fig. 5 (C)

where $m(S_E^u)$ indicates the materials the respective subdomains are part of, and the term $|\cdot| > 0$ indicates that the subphases need to connect via an interface of non-zero length/area. The operator ∂_n^k denotes the k -th normal derivative. The jump in normal derivatives operates on the extension of the test and trial functions $\tilde{\mathbf{v}}^h$ and $\tilde{\mathbf{u}}^h$ from their respective subdomains S_{E+}^u and S_{E-}^u to the background elements Ω^{E+} and Ω^{E-} which are obtained by omitting the enrichment function ψ_B^e in (11). For details on the scaling of the penalty parameter γ_G^k , the reader is referred to [62].

Using the definition of clusters introduced in Sect. 2.2, the goal of the implementation of the ghost stabilization in the preprocessor consists of finding and constructing the cluster pairs corresponding to each ghost facet and valid subdomain combination. For the example shown in Fig. 5 (C), we need to construct the cluster pairs shown in Fig. 6.

3 Mesh generation

This section outlines the implementation of the preprocessor with its associated algorithms and data structures. Furthermore, it discusses the concepts and strategies used to achieve robustness and parallel scalability. The input to the preprocessor consists of one or multiple geometries and a background mesh in which they are immersed. The output are clusters, as introduced in Sect. 2.2. These inputs and outputs are formal-

ized in Sect. 3.1, alongside some basic concepts and notation used throughout this section. The preprocessor’s workflow follows a series of four main steps as visualized in Fig. 7 and later summarized in Algorithm 1. First, a grid conforming to the geometric boundaries referred to as the "foreground mesh" \mathcal{T} , is generated. This is achieved by triangulating the background mesh in two sub-steps: regular and templated subdivision, both outlined in Sect. 3.2. In the second step, topological information needed for the enrichment and ghost stabilization is computed and stored for the foreground mesh \mathcal{T} , see Sect. 3.3. Using the topology information, enrichment is performed in the third step outlined in Sect. 3.4. The generation of ghost clusters is discussed as the last step in Sect. 3.5. The parallel implementation of each of these steps is discussed separately at the end in Sect. 3.6, with the intent of not providing too much information at once.

3.1 Preliminaries

Inputs and Outputs

The goal of the immersed FE preprocessor is to generate sets of clusters associated with the various subdomains, namely, material domains Ω_m , boundaries Γ_m , interfaces $\Gamma_{m,l}^I$, and ghost faces \mathcal{F}_G , to support immersed FE analysis. This information is formalized in the data structure shown in Fig. 8 containing the background element, which itself consists of geometry (vertices) information and a finite element basis

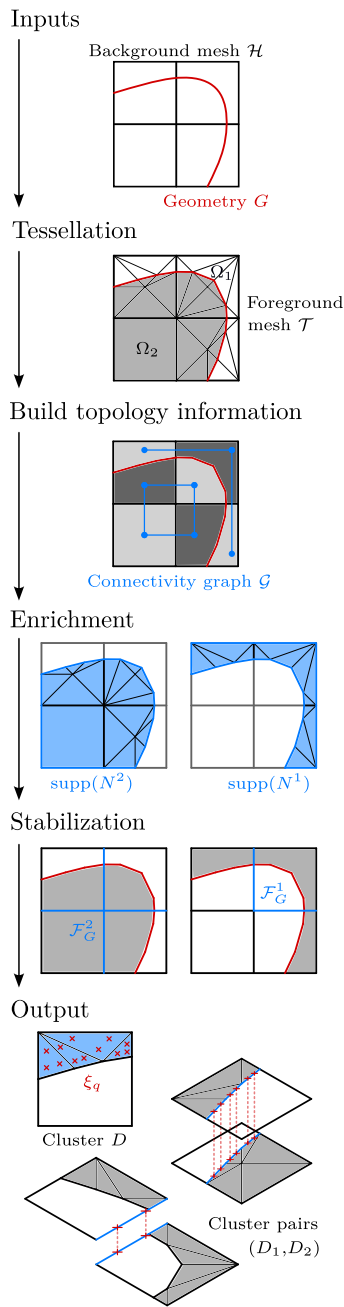


Fig. 7 Sketch of the workflow used in the implementation of the pre-processor

\mathcal{B}_E , and the custom quadrature rules needed in the immersed setting. For internal interfaces $\Gamma_{m,l}^I$ and ghost faces \mathcal{F}_G , clusters are stored in pairs (D_1, D_2) , with matching local ordering of the quadrature points to compute the integrals (16). The local-to-global basis function map on each background element, or the "IEN",² adopting the language in [42], provides

² The naming of the IEN in [42] stems from the fact that it is used to "index elemental nodes". However, its usage is equally valid for non-nodal finite element bases.

Algorithm 1 Driver algorithm.

Input: Background mesh \mathcal{H} , Geometries $\{G_g\}_{g=1}^{n_g}$
 Output: Sets of bulk \mathcal{D}_B and side clusters \mathcal{D}_S ; interface \mathcal{D}_I and ghost side cluster pairs \mathcal{D}_G

```

1: # Foreground mesh generation                                ▷ see Sect. 3.2
2:  $\mathcal{T} \leftarrow \text{initialize\_foreground\_mesh}(\mathcal{H})$ 
3:  $\mathcal{T}.\text{regular\_subdivision}(\{G_g\}_{g=1}^{n_g})$                 ▷ see Algorithm 2 in
   Sect. 3.2.1
4: for each geometry  $G \in \{G_g\}_{g=1}^{n_g}$  do
5:    $\mathcal{T}.\text{templated\_subdivision}(G)$                         ▷ see Algorithm 3 in
   Sect. 3.2.2
6: end for
7: assign parallel IDs to the foreground cells
    $\mathcal{T}.\text{communicate\_cell\_IDs}()$ 
8: for each fg. cell  $c \in \mathcal{T}.C$  do
9:    $c.m \leftarrow \text{material\_map}(c.m)$ 
10: end for
11:
12: # Generate topology information                            ▷ see Sect. 3.3
13:  $\mathcal{T}.\mathcal{F} \leftarrow \text{compute\_entity\_connectivity}(\mathcal{T}, d - 1)$  ▷ see
   Algorithm 12
14: fg. facets attached to each bg. facet DF  $\leftarrow$ 
    $\mathcal{T}.\text{compute\_bg\_facet\_descendants}()$                     ▷ see Algorithm 4
15:  $\mathcal{T}.\text{generate\_subphases}()$                                ▷ see Algorithm 5
16:  $\mathcal{T}.\text{generate\_subphase\_graphs}(\text{DF})$                    ▷ see Algorithm 6
17:
18: # Enrichment - create unzipped background elements      ▷ see Sect.
   3.4
19:  $\mathcal{E} \leftarrow \text{unzip\_interpolation\_mesh}(\mathcal{H}, \mathcal{T})$     ▷ see
   Algorithm 7
20:
21: # Cluster generation
22:  $\mathcal{D}_B \leftarrow \text{create\_bulk\_clusters}(\mathcal{T}, \mathcal{E})$         ▷ see Algorithm 8
23:  $\mathcal{D}_S, \mathcal{D}_I \leftarrow \text{create\_side\_clusters}(\mathcal{T}, \mathcal{E})$  ▷ see Algorithm
   9
24:  $\mathcal{D}_G \leftarrow \text{generate\_ghost\_clusters}(\mathcal{T}, \mathcal{E})$  ▷ see Algorithm
   10 in Sect. 3.5
25:
26: return  $\mathcal{D}_B, \mathcal{D}_S, \mathcal{D}_I, \mathcal{D}_G$ 
    
```

the information necessary to assemble the elemental residual and tangent matrices into a global system.

The inputs needed for constructing clusters are a background mesh \mathcal{H} and a set of geometries $\{G_g\}_{g=1}^{n_g}$ defining the various material domains Ω_m . The background mesh is a finite element mesh that provides: (i) a list of elements that make up the support of a basis function with global index B and (ii) the connectivity information of its mesh entities as discussed in the following paragraph. Without loss of generality, we assume a rectilinear background mesh of Quad or Hex elements, but the algorithm also applies to unstructured meshes made of different element types. Non-conformal meshes, such as those found in locally refined hierarchical meshes, require additional modifications, which we may address in future work. The representation of each geometry needs to support the following queries:

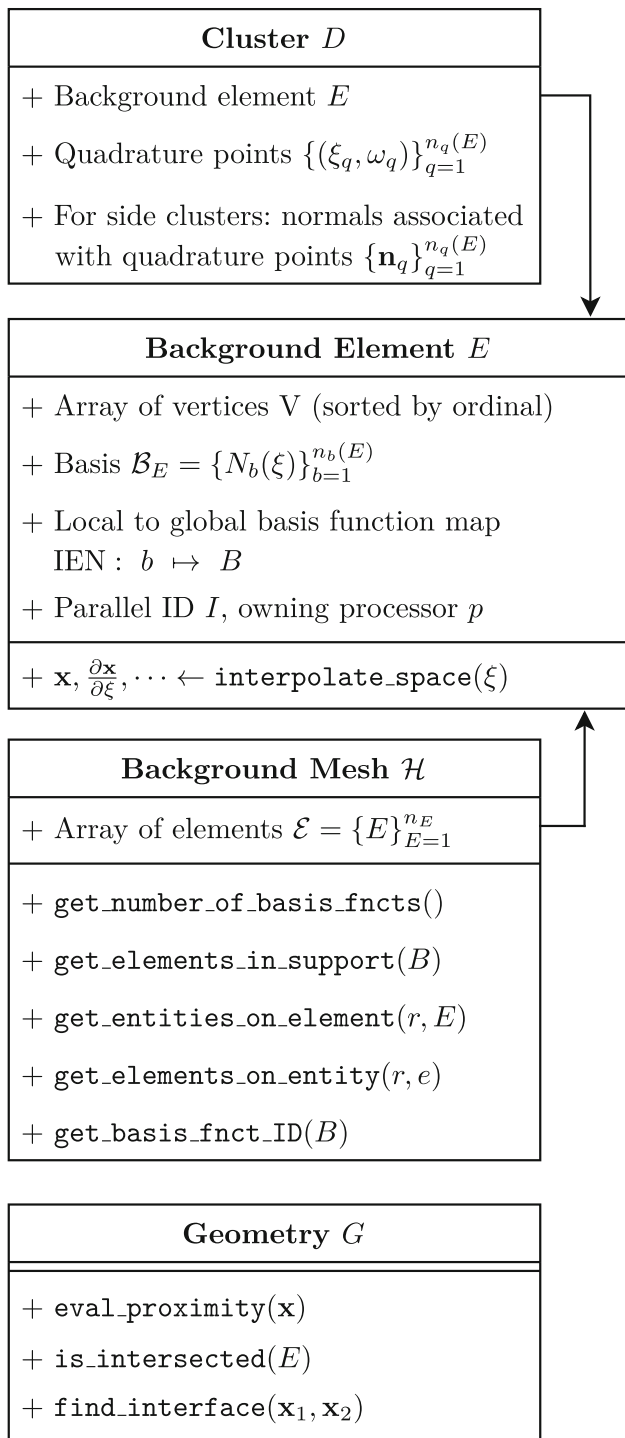


Fig. 8 Unified modeling language (UML) chart outlining the data structures for the input and output data for the preprocessor




Fig. 9 Entity connectivity consisting of two maps

tries can be translated into implicit geometries [84], or approaches such as ray-tracing [4] may be used.

- Check whether a given background element E is intersected by a geometry interface. Note that, for robustness, the query may not return any false negatives, though false positives are acceptable and only affect computational efficiency. For implicit geometries, the range evaluation procedure presented by Saye [70] may be employed.
- Find the location of an interface between two points with opposite proximity within a background element. For implicit geometries, this can be done by interpolating the level-set field along the edge or on the background mesh and then performing a polynomial root-finding operation [56, 66, 70]. For geometries defined through boundary representations, approaches such as the one in [8] may be utilized.

Mesh Entities and Connectivity

Understanding the relationships between mesh entities – vertices, edges, faces (2D cells), and 3D cells – is essential for the algorithms presented in this paper. The following mesh entity and connectivity terminology is used throughout this paper:

- **Rank r :** The dimensionality of a mesh entity, e.g., vertices have rank 0, edges rank 1, faces rank 2, and cells rank 3.
- **Facet F :** A $(d-1)$ -dimensional entity of a d -dimensional mesh, e.g., edges in 2D and faces in 3D.
- **Ordinal o :** The local index of an entity with respect to another entity, as shown in Fig. 10 (A).
- **Cells c and Vertices v :** Purely geometric entities, distinct from elements and nodes, which also include finite element basis information. Vertices specifically refer to the corner points of a cell.
- **Entity Index:** A (processor-local) identifier for mesh entities of a specific type, ranging from 1 to the total number of such entities on the processor.

Background meshes generated with commonly used mesh generators do not provide this information directly, but it can be computed using Algorithm 12 provided in the appendix.

Using these definitions, we introduce two maps, implemented as nested arrays, to express the connectivity of mesh

- Evaluate the proximity of a point in space, i.e., whether a point lies outside, inside, or on the boundary of the geometry within some tolerance. Implicit geometry definitions naturally lend themselves to this query. Explicit geome-




Fig. 10 **A** Vertex and edge ordinals for a triangular cell. **B** Vertex, edge, and cell indices for an example mesh. **C** Cell-vertex connectivity C defining the mesh in **(B)** and the resulting edge-to-cell (EtC) and cell-to-edge (CtE) connectivity for the mesh in **(B)**. A detailed annotation is found in Sect. B.2

entities: (i) an entity-to-cell map (EtC), which contains the cells attached to an entity of a given index, and (ii) a cell-to-entity map (CtE), which contains the indices of entities attached to a cell of a given index. The list of entities attached to a cell is ordered by the ordinal of the entity with respect to the cell, as illustrated in Fig. 9. An example of these maps is shown in Fig. 10 (C) for the edge connectivity of the small four-element mesh shown in Fig. 10 (B), with a more complete explanation provided in Sect. B.2.

Notation and Conventions

The following notation and conventions are used in the presented algorithms.

- The "." operator denotes access to a member function or variable of an object.
- Arrays are represented using braces "{}" and tuples using parentheses "()".
- The presented algorithms use 1-based indexing unless otherwise specified, for the sake of readability.
- Access to an element of an array or a map is denoted using the "[]" operator.
- "[-1]" is shorthand for accessing the last element of an array.
- Loops over each element of an array are condensed using the shorthand "[:]".
- All relevant objects have an index. The symbols representing an object are used interchangeably as the index

or the object itself; e.g., in " $E \in \mathcal{H.E}$ ", E may represent the element or its index in the global list $\mathcal{H.E}$.

- Lines in italics and starting with a "#", are comments.
- Components of algorithms are colored in gray to indicate they are only relevant to parallel implementation. We only discuss these additions at the end of Sect. 3.6 and recommend ignoring them on the first read-through.
- The words foreground and background are abbreviated as "fg." and "bg.", respectively.

Using these conventions, the workflow is summarized in Algorithm 1.

3.2 Conformal foreground mesh generation

The goal of this part of the preprocessor algorithm is to subdivide the background mesh's elements intersected by geometric interfaces into groups of triangular or tetrahedral cells with facets conforming to those interfaces. As a result, the geometric interfaces are approximated by sets of straight edges in 2D or flat triangles in 3D, while the volumetric domains for each of the materials Ω_m are approximated by groups of the generated triangular or tetrahedral cells. The focus for this step, besides enabling the generation of topology information for enrichment and ghost stabilization, is on robustness and speed. To this end, a combination of templated subdivision procedures is employed.

A drawback of the presented approach is that the quadrature rules, which are subsequently constructed from the triangulation in Sect. 3.4, only capture the geometry and its interfaces with linear accuracy. However, modifications using isoparametric projection [47] or manipulation [76, 77] of the resulting triangulated foreground mesh are possible as part of the foreground mesh generation; although, they are not further addressed in this work.

To obtain a Tri/Tet mesh in the vicinity of the interfaces, background elements cut by any geometric interfaces are first subdivided into a pre-defined set of triangular or tetrahedral cells. This step is referred to as the "regular subdivision". On the resulting cells, the edges are probed for intersections with geometric interfaces and further subdivided into triangular or tetrahedral cells conforming to the interface using a set of subdivision templates. This "templated subdivision" approach has been used commonly in other works, see [82]. The main difference in the presented approach is the addition of background ancestry information, which enables the computation of topology information without floating point operations at a later stage. Lastly, a recursive procedure to determine the material membership of each generated cell is incorporated into the subdivision processes. The complete process is illustrated in Fig. 11.

In what follows, we first introduce the concept of background ancestry and the initialization of the foreground mesh,

before discussing the regular subdivision, the templated subdivision, and material assignment.

Background Ancestry

The background ancestor of an entity of the foreground mesh is the background entity with index a of the lowest rank r , i.e., dimensionality, which fully contains the foreground entity. This information, stored in the form of index pairs (r, a) , is used to unambiguously identify vertices during the subdivision and to generate topology information without the use of floating point arithmetic. Examples of the background ancestry of various entities formed during the conformal mesh generation process are shown in Fig. 12.

Consider the tetrahedral foreground cell $\{4, 8, 9, 15\}$ in Fig. 12 that was formed during the subdivision process. Vertices 4 and 8 were part of the original background mesh and are descendants of those vertices; hence, their background ancestry is denoted as $(0, 4)$ and $(0, 8)$ for background vertices (rank $r = 0$) with indices $a = 4$ and $a = 8$, respectively. Vertices 9 and 15 were formed on a face of the background element; hence, their ancestors are denoted as $(2, 4)$ and $(2, 3)$, for being descendants for background faces (dimensionality, i.e. rank $r = 2$) with indices $a = 4$ and $a = 3$. A new vertex v_3 is formed on the edge connecting vertices 4 and 8. It can be checked using entity connectivity information whether those vertices are attached to the same edge or face, or are part of the same background element. In this case, they are part of the same background edge ($r = 1, a = 12$), and therefore any new vertex formed by them would be a descendant of that edge. Considering entities of one dimension higher, the face formed between the vertices v_1, v_2 , and v_3 would be a descendant of the background element (dimensionality $r = 3$) itself, as two of the vertices descend from different faces.

Initialization

To initialize the foreground mesh data, shown in Fig. 13, the geometric information, i.e., the vertex coordinates and vertex connectivity of the background elements, is copied from the background to the foreground mesh. The ancestry (r, a) for every foreground vertex and cell E is initialized trivially, given that each foreground vertex or cell is a descendant of the background vertex or element with the same index. The material index $m = 0$ is assigned to all cells.

To access foreground cell information from the perspective of individual background elements, child mesh containers CM are created, containing the vertices and cells formed within each background element. This includes the parametric coordinates of the vertices relative to the parametric coordinate space of the background element, which are essential for geometry interpolation operations. During initialization, each child mesh container is populated with the vertices from the single cell that originally constitutes the background element.

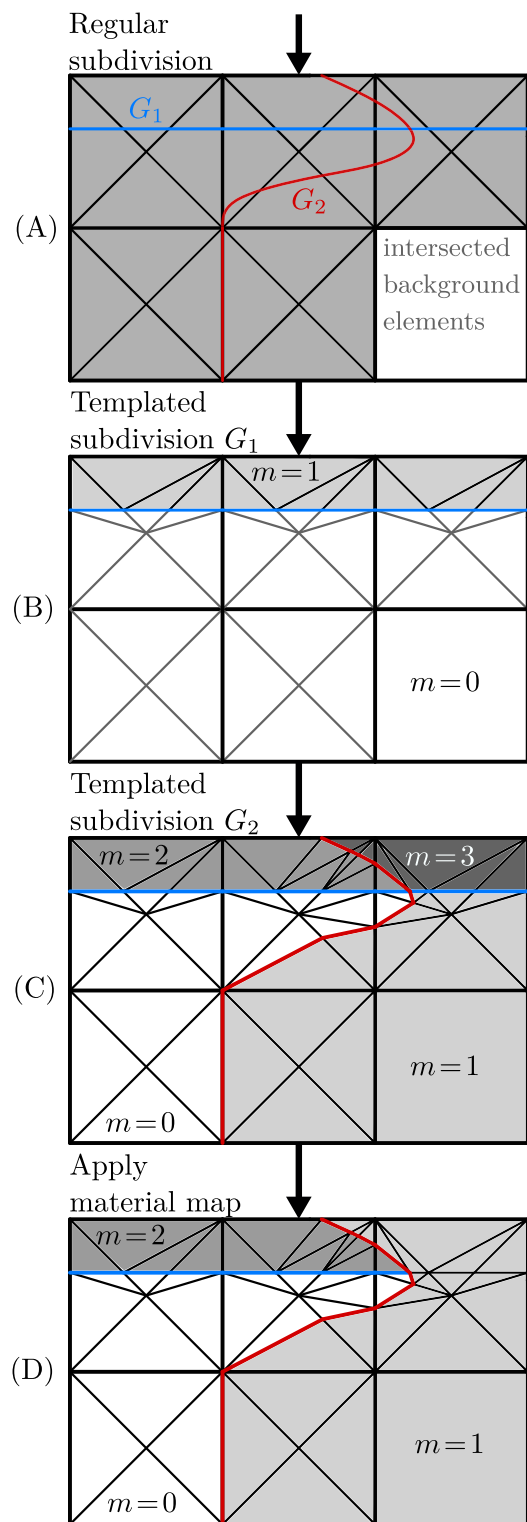


Fig. 11 Subdivision process overview. **A** Regular subdivision: application of templates to bg. elements intersected by geometric interfaces, shown with their original accuracy. **B** and **C** Templated subdivision with recursive material assignment for the first G_1 and second geometry G_2 , respectively. **D** Application of a material map and the resulting foreground mesh




Fig. 12 **A** Concept of background ancestry demonstrated using a Tet which gets subdivided by an interface. The "parent" background element is indicated in dashed lines. The ancestry (entity rank r , entity index a) of the existing (blue) and added (red) vertices are shown alongside some of the newly formed faces (gray). **B** Indices of the edges (green) and faces (red) attached to the parent background element

The remaining information in Fig. 13 is related to the material topology and is populated after the subdivision process.

3.2.1 Regular subdivision

The regular subdivision serves the purpose of generating Tri or Tet cells near the geometric interfaces as these allow for a robust conformal mesh generation procedure using templates. If the algorithm is to be applied to a pure Tri/Tet background mesh, this step is skipped. For the procedure outlined in Algorithm 2, the templates shown in Fig. 14 are applied to every background element intersected by a geometric interface.

The templates employed here do not minimize the number of Tris or Tets required to subdivide a Quad or Hex element. However, they offer two distinct advantages over templates generating a minimal number of Tri/Tet cells. First,

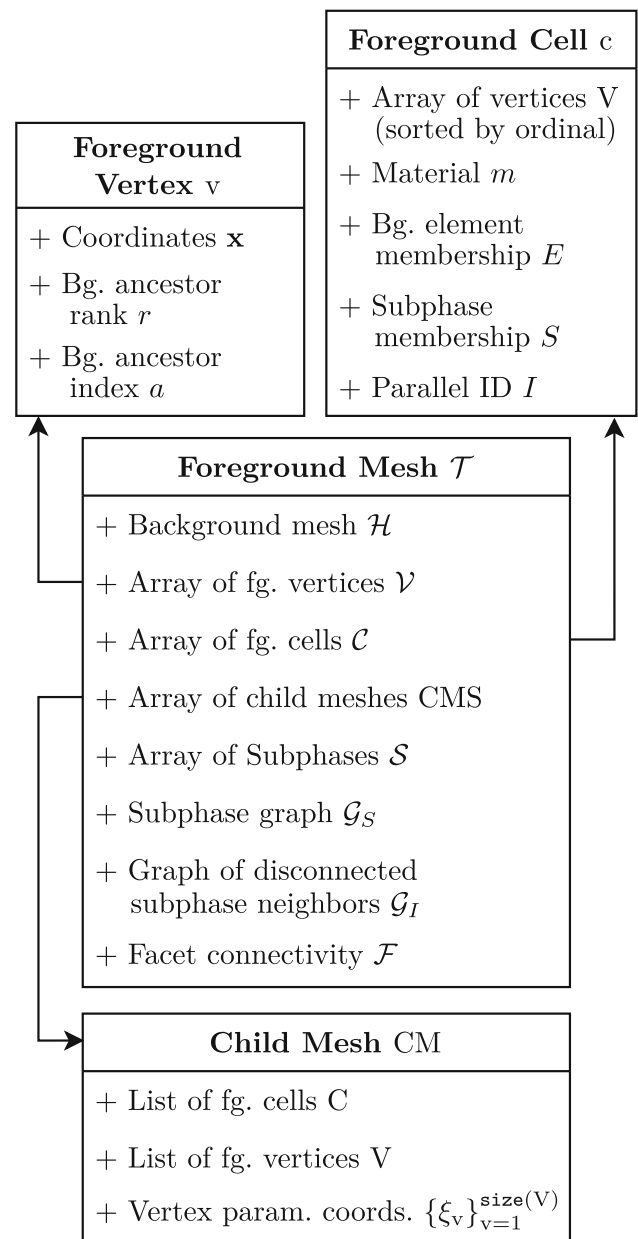


Fig. 13 UML chart showing the data stored in the foreground mesh. The first four members listed for the foreground mesh are populated during the subdivision process in Sect. 3.2; the remaining four are filled during the generation of the topology information in Sect. 3.3

the cells generated by these templates do not further subdivide existing background edges, ensuring that each generated foreground cell includes at most one complete background edge in both 2D and 3D. This simplifies the determination of the background ancestry of newly formed vertices and facets during the templated subdivision, reducing the process to a series of straightforward checks outlined in Sect. B.1. Second, the templates exhibit symmetry about all axes, which ensures that, for higher-order geometry descriptions, the resulting lower-order approximation does not introduce

a bias depending on the orientation of the applied template.

To perform the regular subdivision, as outlined in Algorithm 2, the templates contain the following information:

1. A list of new vertices to be generated with their location ξ in the parametric space of the background element in addition to the rank r and ordinal o of the background entity they are located on. This information is provided through the function `get_new_vertex_locations()`.
2. A pre-defined cell-vertex connectivity for the new cells in the format shown in Fig. 10. The template expresses the connectivity using the indices of both the existing and new vertices. This information is provided through the function `generate_cells(E.V, NVI)`, where $E.V$ is the list of vertices of the existing cell and NVI is the list of indices of the new vertices generated by the template.

New vertices and cells are temporarily stored in a queue, rather than directly added into the foreground mesh data structure. A queue consists of an array of requested entities with their initialization data and a (hash) map, which relates identifying information with the position of the entity in the array of requests. This approach

1. avoids repeated resizing operations of the global lists of vertices \mathcal{V} and cells \mathcal{C} on the foreground mesh, which can be computationally costly,³
2. facilitates the generation of unique vertices. The regular subdivision templates, shown in Fig. 14, generate at most a single vertex on a given background entity. The rank and index of the background entity are used to identify identical vertices that may have already been requested by a neighboring background element if the new vertex is located on a shared background facet. Otherwise, newly created vertices would need to be merged using their physical coordinates, as is done by [93], requiring floating point arithmetic.

Lastly, to eliminate the subdivided cell from the foreground mesh, its vertex list is overwritten with the vertex list of one of the new cells. This avoids a delete operation and the need for re-indexing.

³ Dynamic arrays, such as C++’s `std::vector` are equipped with smart resizing strategies when repeatedly appending elements to the array. Such strategies may lead to over-allocation of memory and repeated copying of the existing array elements, resulting in significant computational cost and memory fragmentation in the context of the large arrays being edited in this case.




Fig. 14 Regular subdivision templates in 2D (A) and 3D (B). New vertices generated by the templates are shown with their ancestry (ancestor entity rank, ordinal) in red and existing vertices in blue

3.2.2 Templated subdivision

The templated subdivision completes the generation of the foreground mesh by subdividing the cells generated in the regular subdivision into cells conforming to the geometric interfaces. This procedure is outlined in Algorithm 3.

The procedure assesses foreground edges generated during the regular subdivision if they are intersected by a geometric interface and then generates additional vertices and cells according to one of the templates shown in Fig. 15 or their permutations. Depending on the proximity of vertices, a new material index m is assigned to the elements. The process is repeated for every geometry applied.

First, the proximity of each foreground vertex, i.e., whether a vertex is located inside, outside, or on the boundary of a geometry, is determined and stored to avoid error-prone re-computation. Subsequently, the edges of every potentially intersected cell are checked for differences in the proximity of the adjacent two vertices. If so, it is considered intersected and the intersection point is determined. The approach does

Algorithm 2 Regular Subdivision.

Input: Initialized foreground mesh \mathcal{T} (self), Geometries $\{G_g\}_{g=1}^{n_g}$

Output: Foreground mesh \mathcal{T} with triangulated cut bg. elements

```

1: initialize new vertex  $Q_v$  and cell queues  $Q_c$ , get template  $T \leftarrow$ 
   regular_subdivison_template( $d$ )
2: get location of new vertices in bg. element  $\{(\xi, r, o)\}_{i=1}^{\text{NumNewVerts}} \leftarrow$ 
    $T.get\_new\_vertex\_locations()$ 
3: for each bg. element  $E \in \mathcal{H.E}$  do
4:   if  $G$ .is_element_intersected( $E$ ) for any  $G \in \{G_g\}_{g=1}^{n_g}$ 
   then
5:     for each new vertex in template  $(\xi, r, o) \in$ 
    $\{(\xi, r, o)\}_{i=1}^{\text{NumNewVerts}}$  do
6:       get the entity index  $e =$ 
    $\mathcal{H}.get\_entities\_on\_element(E, r)[o]$ 
7:       if this is a new vertex ! $Q_v.request\_exists(r, e)$ 
       then
8:         compute the physical coordinates  $\mathbf{x} \leftarrow$ 
    $E.interpolate\_space(\xi)$ 
9:         request new vertex and remember its index
    $NVI.append(Q_v.queue(\mathbf{x}, r, e))$ 
10:        else # another bg. element has already requested this vertex
11:          store the index for the new vertex
    $NVI.append(Q_v.get\_index(r, e))$ 
12:        end if
13:        store the new vertex in the child mesh
    $CMS[E].add\_vertex(NVI[-1], \xi)$ 
14:        end for
15:        get the cell-vertex connectivity for the new cells  $C_{new} \leftarrow$ 
    $T.generate\_cells(E.V, NVI)$ 
16:        for each new cell in template except the first  $c_{new} \in$ 
    $C_{new}[-1 : 2]$  do
17:          queue cell and add its index to the child mesh
    $CMS[E].C.append(Q_c.queue(c_{new}, E, 0))$ 
18:        end for
19:        replace the subdivided element with the first new cell
    $C[E].V \leftarrow C_{new}[1]$ 
20:        end if
21:      end for
22: create_vertices( $Q_v$ )
23: create_cells( $Q_c$ )

```

not allow for geometry intersections on edges that contain a vertex that lies on the interface. It also limits the number of intersections to one per edge. This ensures robustness toward unwanted intersection configurations but leads to a dependency of the final foreground mesh on the order in which the geometries are applied. Consider G_2 in Fig. 11 (A): the protrusion on the right would be cut off if G_2 were to be applied first, as it intersects the background edge twice. The process would not consider the edge intersected in this case as both its vertices, i.e., its end-points, lie on the same side of the interface. The issue of insufficiently resolved geometry can be overcome by local mesh refinement, either by refining the background mesh or by applying additional, e.g., quadtree/octree subdivision templates before the regular subdivision.

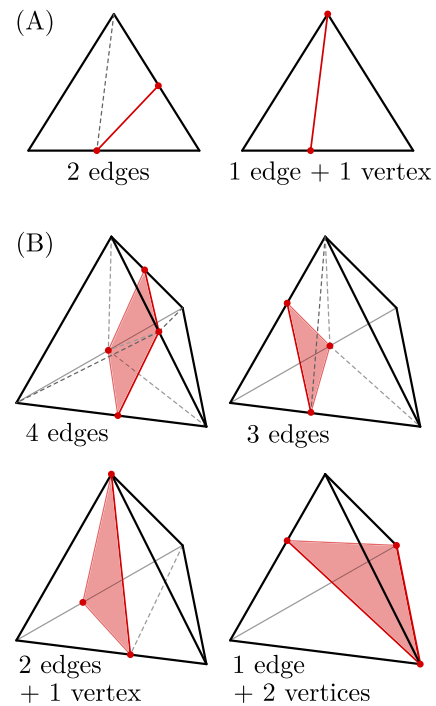


Fig. 15 Possible intersection configurations for foreground cells during the templated subdivision in 2D (A) and 3D (B). The templates are permutations of the configurations shown

Computing the interface location along the edge, both the physical and parametric coordinates of the new vertex are found. The background ancestry of the new vertex is determined by finding a common ancestor of the two edge vertices. For the regular subdivision template presented here, determining the ancestor is simplified as outlined in Appendix B.1. If a different regular subdivision template or an unstructured Tri or Tet input background mesh were to be used, search operations would need to be performed on the entity connectivity of the background mesh to determine the background ancestors of the new vertices. Requests for the new vertices on the intersected edges are queued in the same manner as for the regular subdivision. A sorted pair of vertex indices serves as the identifier for the edge. The cell-vertex connectivity for the new cells is generated by providing the indices of the new and existing vertices to the template through the method `generate_cells($c.V, V_{new}$)` in Algorithm 3, and, again, one of the new cells is used to replace the information of the subdivided cell to forego an explicit deletion.

Material assignment

The material membership of all cells is determined concurrently with the templated subdivision in a recursive process. We have found that performing material assignment after the subdivision process leads to robustness issues, due to the mismatch of discretized interfaces and the geometries of the LS isocontours.

The sign of the proximity of each cell is stored for every geometry applied. Interpreting the resulting Boolean array as an integer, as shown in Fig. 16, yields a material index. Alternatively, this can be stated as a recursive rule applied during the subdivision process:

$$m = 2 \cdot m + \begin{cases} 1, & \text{if } P > 0, \\ 0, & \text{otherwise,} \end{cases} \quad (21)$$

where P represents the proximity of the cell to the geometry.

Given the restrictions on the possible intersection configurations for the subdivision templates, each of the cells generated by a template can necessarily only have vertices attached to it with non-conflicting proximity values, i.e., 0 and +, or 0 and −, but never + and − or only 0’s. Ignoring vertices with 0, the proximity values of the remaining vertices uniquely determine the proximity of the cell. Figure 16 demonstrates this logic for a single subdivided cell in 2D. The rule (21) is also applied to the cells without edges intersected by the geometric interface within a given subdivision step.

In the last step, a user-defined material map may be applied to the foreground mesh to merge different material regions, as shown in Fig. 11 (D). The material map $H : \{0, 1, \dots, 2^{n_g} - 1\} \mapsto \mathcal{M}$ assigns a material index m to each of the 2^{n_g} material domains created during the subdivision process. For the particular example shown in Fig. 11 (D), this material map, if represented as a lookup table, reads as

$$H = \{0, 1, 2, 2\}, \quad (22)$$

where each material region is reassigned the index $m = H(m)$. The application of a material map avoids the need to construct interface conditions between regions associated with the same material and physics but separated by an interface of one of the input geometries. It also allows for the exact representation of sharp corners and edges in the interior of background elements, e.g., when defining a square domain using four planes and merging the outside regions.

3.3 Generation of topological information

As discussed in Sect. 2.2, the goal is to evaluate the integrals (14)-(16). To this end, the disconnected subdomains $\{S_E^u\}_{u=1}^{n_u(E)}$ within a given background element E need to

Algorithm 3 Templated subdivision.

Input: Foreground mesh \mathcal{T} (self), Geometry G

Output: Fully subdivided foreground mesh \mathcal{T} (self)

```

1: evaluate each vertex's proximity P[:] ←
   G.compute_proximity(V[:].x) ▷ possible values (+, 0, −)
2: initialize new vertex Qv and cell queues Qc
3: for each intersected bg. element E ∈ {1, ⋯, nE} do
4:   if bg. element is intersected G.is_intersected(E) then
5:     for each child cell in bg. element c ∈ CMS[E].C do
6:       initialize array with empty entry for every edge on cell
       Vnew ← {0}e=1NumEdges
7:       for each edge on cell (v1, v2)e ∈ {(v1, v2)e}e=1NumEdges do
8:         if edge is intersected P[v1] ≠ P[v2] ≠ 0 then
9:           get the location of the interface on edge ξe ←
             G.find_interface(v1.x, v2.x)
10:          ξ ← interpolate_on_edge
11:          (ξe, CMS[E].get_coords(v1),
12:           CMS[E].get_coords(v2))
13:          if existing vertex Qv.request_exists((v1, v2))
14:            then
15:              Vnew[e] ← Qv.get_index((v1, v2))
16:            else
17:              get the physical coordinates x ←
                 E.interpolate_space(ξ)
18:              compute the ancestry (r, a) ←
                 find_common_ancestor(v1, v2)
19:              Vnew[e] ← Qv.queue((v1, v2), x, r, a)
20:            end if
21:          CMS[E].add_vertex_if_not_found(Vnew[e])
22:          end if
23:          end for
24:          copy the proximity of vertices on c into an array Pc[:] =
             P[c.V[:]]
25:          select template using intersected edges & proximity T ←
             select_template(Vnew ≠ 0, Pc, c)
26:          for cell-vertex connectivity for each new cell cnew ∈
             T.generate_cells(c.V, Vnew) do
27:            determine new cell's proximity and material m ← 2 ·
             c.m + is_positive(vote(P[cnew.V[:]]))
28:            (# for vote function, ignore new vertices which are
             outside the index range of P)
29:            if iteration is not the last in loop then
30:              CMS[E].V.append(Qc.queue(cnew, E, m))
31:            else # replace information for subdivided cell with one
             of the new cells
32:              c.V ← cnew, c.m ← m
33:            end if
34:          end for
35:        end for
36:      else # non-cut bg. elements
37:        determine proximity of the cell P ← vote(P[E.V[:]])
38:        for each of the child cells c ∈ CMS[E].C do
39:          increment the material index c.m ← 2 · c.m +
             is_positive(P)
40:        end for
41:      end if
42:    end for
43:  create_vertices(Qv)
44:  create_cells(Qc)

```




Fig. 16 Recursive material assignment rule (21) applied during the templated subdivision. The proximity P of each cell after subdivision is obtained by applying a vote to their attached vertices. The angle brackets $\langle \cdot, \dots \rangle$ indicate arrays of Booleans, which are interpreted as binary numbers to obtain material index m

Subphase S
+ Bg. element membership E
+ Ordinal u
+ Array of fg. cells C
+ Parallel ID I

Fig. 17 Data structure for the subphase




Fig. 18 Example for the foreground facets descending from two background facets in 3D. On the left: non-conformal transition between a non-cut and a cut element due to the regular subdivision template. On the right: conformal transition with an intersecting interface

be identified. From here on, the term "subphase" is used to refer to these disconnected subdomains. In code, a subphase is represented by a list of foreground cells that constitute a connected material subdomain within a background element. The associated data structure is shown in Fig. 17.

Following the discussion in Sects. 2.1 and 2.3, connected material regions within each basis function's support need

Algorithm 4 Generate descendants for the interior bg. facets
 Input: fg. mesh \mathcal{T} (self) with facet connectivity \mathcal{F}
 Output: fg. mesh \mathcal{T} (self), descendant fg. facets of each bg. facet DF

```

1: initialize list of descendant facets for each bg. facet DF =
   {}f=1H.get_num_entities(d-1)
2: for each fg. facet  $F_{fg} \in \mathcal{F}.CtE[c]$  on each fg. cell  $c \in \mathcal{T}.C$  do
3:   if  $F_{fg}$  has not already been processed then
4:     get the vertices attached to the current facet  $V_F \leftarrow$ 
       c.get_entity(d-1,  $F_{fg}$ )
5:     bg. elements vertices are contained in A  $\leftarrow$ 
        $\bigcap_{i=1}^{size(V_F)} \mathcal{H}.get\_entities\_on\_cell(V_f[i].r, V_f[i].a)$ 
6:     if size(A) = 2, i.e., the fg. facet is contained on a bg. facet
       then
7:       get the bg. facet  $F_{bg} \leftarrow$ 
            $\bigcap_{i=1}^2 \mathcal{H}.get\_entities\_on\_cell(d-1, A[i])$ 
8:       store that the current fg. facet coincides with bg. facet
           DF[ $F_{bg}$ ].append( $F_{fg}$ )
9:     end if
10:  end if
11: end for
12: return DF
    
```

to be identified for the enrichment and ghost stabilization strategies. Additionally, an easy identification of interfaces between each of the material regions is useful. The pre-processing steps in this subsection generate, besides the subphases themselves, two graphs: one to represent the connection between the subphases \mathcal{G}_S and one to represent adjacent but disconnected material regions \mathcal{G}_I . Both graphs are stored in the form of adjacency lists.

Generation of Facet Connectivity Information

A connectivity for the facets on the foreground mesh is generated using Algorithm 12. The regular subdivision templates in 3D, see Fig. 14 (B), resulting in a foreground mesh with hanging vertices, where subdivided and non-subdivided background elements meet. The left background facet in Fig. 18 depicts such a configuration. Due to the mismatching facets, the connectivity information between the adjacent foreground cells needs to be established separately.

The background ancestry of the foreground vertices and the entity connectivity of the background mesh are used to find the missing connections in Algorithm 4. The procedure is based on the fact that any foreground facet coinciding with an interior background facet is attached to exactly two background elements. First, the background ancestors for each vertex on each foreground facet are found. If the intersection between the lists of cells attached to the respective ancestors returns exactly two background element indices, the foreground facet must be a subset of the background facet connecting the two background elements. The resulting list of descendant foreground facets DF for each interior back-

Algorithm 5 Generate the subphases.

Input: subdivided fg. mesh \mathcal{T} (self)

Output: fg. mesh \mathcal{T} (self) with subphases \mathcal{S}

```

1: initialize interface subphase graph  $\mathcal{G}_I \leftarrow \{\{\}\}$  and list of subphases
    $\mathcal{S} \leftarrow \{\}$ 
2: for each bg. element  $E \in \mathcal{H.E}$  do
3:   generate graph of elements in child mesh  $\mathcal{G}_{local} \leftarrow$ 
   generate_element_graph(CMS[E].C,  $\mathcal{F}$ )
4:   get the materials for each of the fg. cells  $M[:]\leftarrow$  CMS[E].C[:].m
5:   get groups of connected fg. cells within bg. element  $\{C_u\}_{u=1}^{n_u(E)} \leftarrow$ 
   flood_fill( $\mathcal{G}_{local}$ , M)
6:   for each connected group of cells  $C_u \in \{C_u\}_{u=1}^{n_u(E)}$  do
7:     create a new subphase
8:      $\mathcal{S}.append(create\_subphase(C_u, E, u))$ 
9:     for each fg. cell in group  $c \in C_u$  do
10:      assign the subphase index to the cell  $c.S \leftarrow size(\mathcal{S})$ 
11:     end for
12:   end for
13:   for each subphase  $C_u \in \{C_u\}_{u=1}^{n_u(E)}$  do
14:     find subphase indices of adjacent elements  $S_a \leftarrow$ 
     collect_subphase_neighbors( $C_u, \mathcal{G}_{local}, \mathcal{F}$ )
15:     add the neighbors to the subphase interface graph
      $\mathcal{G}_I.append(S_a)$ 
16:   end for
17: end for
18: communicate_subphase_IDs()

```

ground facet is used to establish the subphase connectivity across background elements.

Identification and Generation of Subphases

The subphases themselves are created by component analysis. A cell connectivity graph \mathcal{G}_{local} is generated on each child mesh and followed by a flood-fill operation considering their material membership, as stated in Algorithm 5. The previously generated facet connectivity \mathcal{F} lends itself to constructing the cell connectivity, as it represents its dual-graph. Subphases are initialized from each of the containers of foreground cells returned by the flood-fill. In addition, the index of the subphases is stored with the cells. The subphases resulting from performing these operations on the foreground mesh in Fig. 11 (D) are visualized in Fig. 19 (A). Further, the local connectivity information already constructed is used to identify neighboring subphases and store them in the interface graph \mathcal{G}_I . As the subphases are disconnected by definition, subphase neighbors inside a background element are never part of the graph \mathcal{G}_S . The initialized graph \mathcal{G}_I is shown in Fig. 19 (B).

Generation of the Subphase Graphs

For the subsequent construction of \mathcal{G}_S and completion of \mathcal{G}_I , only subphase connections across background elements need to be considered. Hence, only foreground facets fully contained on the interior background facets, i.e., those in DF, are needed to find the connections. The procedure outlined in Algorithm 6 iterates over these facets and stores the subphase connections in the adjacency lists \mathcal{G}_S and \mathcal{G}_I .

Algorithm 6 Generate the subphase graphs.

Input: fg. mesh \mathcal{T} (self) with subphases \mathcal{S} , descendant facets of each bg. facet DF

Output: finalized fg. mesh \mathcal{T} (self) with subphase graphs \mathcal{G}_S and \mathcal{G}_I

```

1: initialize empty subphase graph  $\mathcal{G}_S \leftarrow \{\{\}\}_{S=1}^{size(\mathcal{S})}$ 
2: for each bg. facet  $F_{bg} \in \{1, \dots, size(DF)\}$  do
3:   get the bg. elements connected to bg. facet  $\{E_1, E_2\} \leftarrow$ 
    $\mathcal{H}.get\_cell\_on\_entity(d-1, F_{bg})$ 
4:   if facet is on exterior  $E_2 = \text{NULL}$  then
5:     skip it
6:   # if one of the cells is not cut, the material of the bg. facet adjacent
   cells must be the same
7:   else if one of the elements is not cut  $size(CMS[E_1].C) = 1$ 
   OR  $size(CMS[E_2].C) = 1$  then
8:     get the subphases on either side of the bg. facet  $(S_1, S_2)$ 
9:     mark the two subphases as neighbors  $\mathcal{G}_S[S_1].append(S_2)$ ,
      $\mathcal{G}_S[S_2].append(S_1)$ 
10:  else # both adjacent bg. elements are cut
11:    for each fg. facet on bg. facet  $F_{fg} \in DF[F_{bg}]$  do
12:      get the subphases on either side  $\{S_1, S_2\} \leftarrow \mathcal{F}.EtC[:].S$ 
13:      if the neighboring subphases are of the same material
          $S_{1,m} = S_{2,m}$  then
14:         $\mathcal{G}_S[S_i].append\_if\_not\_found(S_j)$ , for  $(i, j) \in$ 
          $\{(1, 2), (2, 1)\}$ 
15:      else
16:         $\mathcal{G}_I[S_i].append\_if\_not\_found(S_j)$ , for  $(i, j) \in$ 
          $\{(1, 2), (2, 1)\}$ 
17:      end if
18:    end for
19:  end if
20: end for

```

Two cases need to be considered for each background facet:

1. One or both background elements are non-cut, as shown on the left of Fig. 18. In this case, any geometric interfaces are necessarily some non-zero distance away from the background facet, and only one subphase is attached to either side of the background facet. The two attached subphases must be of the same material and, hence, are connected. In Algorithm 6, the two subphases are collected from the cells attached to the facets.
2. Both background elements are intersected, as shown on the right of Fig. 18. In this case, the foreground mesh is conformal and, by iterating over the foreground facets, the neighboring subphases are collected and stored. For the special case where a geometric interface coincides with the background facet, the connection is stored in \mathcal{G}_I rather than \mathcal{G}_S .

3.4 Enrichment

To enrich the background basis and construct the enriched basis (13), the supports of the enriched basis functions \tilde{N}_B^e

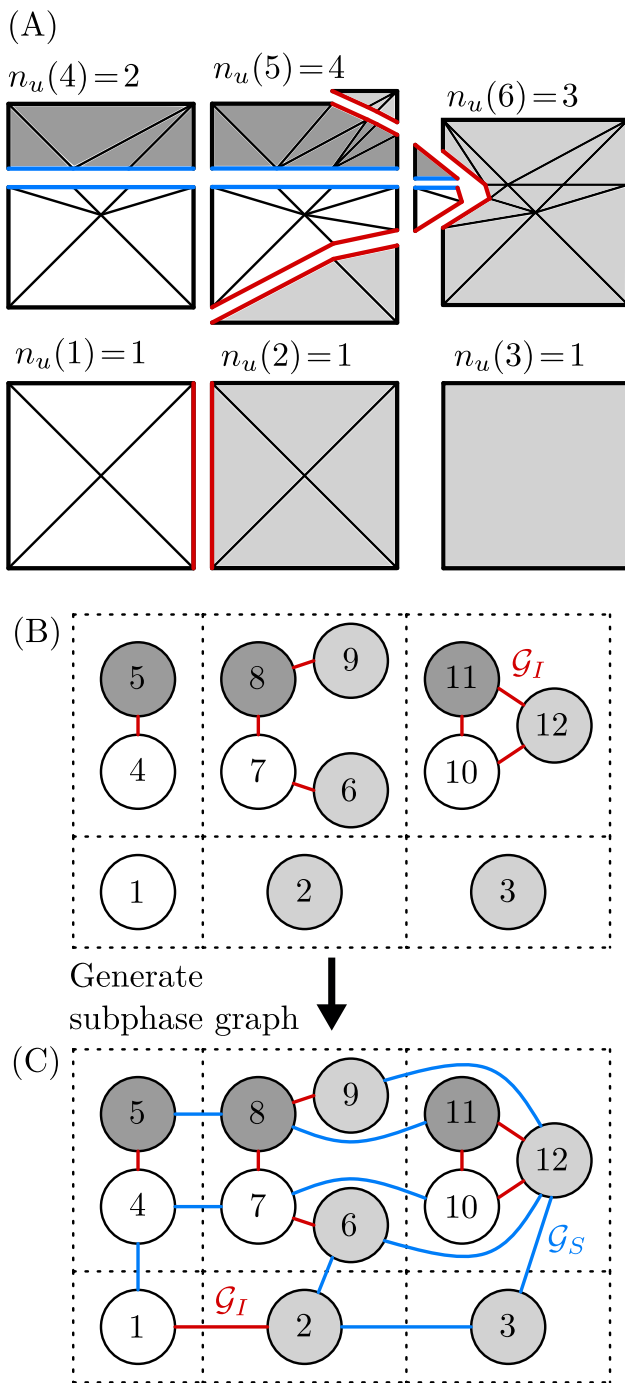


Fig. 19 Generation of subphases and the subphase graph. **A** Identification of disconnected groups of fg. cells in each bg. element using a flood-fill. **B** Subphase interface graph G_I containing neighboring subphases within the background elements. **C** Complete subphase neighbor and interface graphs G_S and G_I

are identified. This requires finding where the enrichment functions $\psi_B^\varepsilon(\mathbf{x})$ are non-zero and the number of "enrichment levels" $n_\varepsilon(B)$ for a given background basis function N_B . The number of enrichment levels $n_\varepsilon(B)$ is computed by performing a flood-fill on the subphase graph within each basis function's support, as outlined in Fig. 20. A list of subphases S_B within the elements comprising $\text{supp}(N_B)$ is collected by pruning the global adjacency list \mathcal{G}_S to $\mathbb{R}_B = \text{supp}(N_B)$. Lists of subphases connected within $\text{supp}(N_B)$ are then obtained by performing a flood-fill algorithm on the pruned graph \mathcal{G}_P . The number of lists is $n_\varepsilon(B)$ and each list corresponds to where $\tilde{N}_B^\varepsilon(\mathbf{x}) \neq 0$.

Unzipping

After identifying the enriched basis functions $\tilde{N}_B^\varepsilon(\mathbf{x})$, the challenge remains to represent and store the resulting enriched basis on the background elements associated with the final cluster output. The strategy chosen here is to represent the enrichment function $\psi_B^\varepsilon(\mathbf{x})$ implicitly by ensuring that a given enriched basis function $\tilde{N}_B^\varepsilon(\mathbf{x})$ is only evaluated at quadrature points $\mathbf{x} \in \text{supp}(\tilde{N}_B^\varepsilon)$, where $N_B(\mathbf{x}) = \tilde{N}_B^\varepsilon(\mathbf{x})$. To achieve this, copies of each of the background elements are created for every subphase within that background element, as shown in Fig. 21. We refer to this process, outlined in Algorithm 7, as "unzipping". The basis functions of the copied background elements $\{\Xi_E^u\}_{u=1}^{n_u(E)}$ will, during the element formation, only be evaluated at points within the associated subphase S_E^u . Hence, the basis functions $\{N_b(\mathbf{x})\}_{b=1}^{n_b}$ on the copied elements do not need to be altered by adding the associated indicator functions ψ_b^ε .

In a next step, a single index ℓ is created to replace the double-index (B, ε) used for the functions $\tilde{N}_B^\varepsilon(\mathbf{x})$. This allows for (11) to be stated as

$$\mathbf{u}^h(\mathbf{x}) = \sum_{\ell=1}^{n_\ell} \tilde{N}_\ell(\mathbf{x}) d_\ell, \tag{23}$$

where n_ℓ is the total number of enriched basis functions $\tilde{N}_\ell(\mathbf{x}) = \tilde{N}_B^\varepsilon(\mathbf{x})$. Subsequently, each of the basis function indices B listed in the IEN of each of the copied background elements Ξ_E^u is replaced with the index of the enriched counterpart $\ell \leftarrow (B, \varepsilon)$ supported within the subphase S_E^u . As a result, the enrichment information is fully encoded in the basis function indices in the IEN-arrays and not further exposed to the outside, so, existing element assembly routines apply to the clusters.

The procedure combining both the enrichment of the basis functions and the unzipping of the background elements is presented in Algorithm 7.

Construction of Clusters

The remaining task is to generate the quadrature points associated with each subphase and package the information up

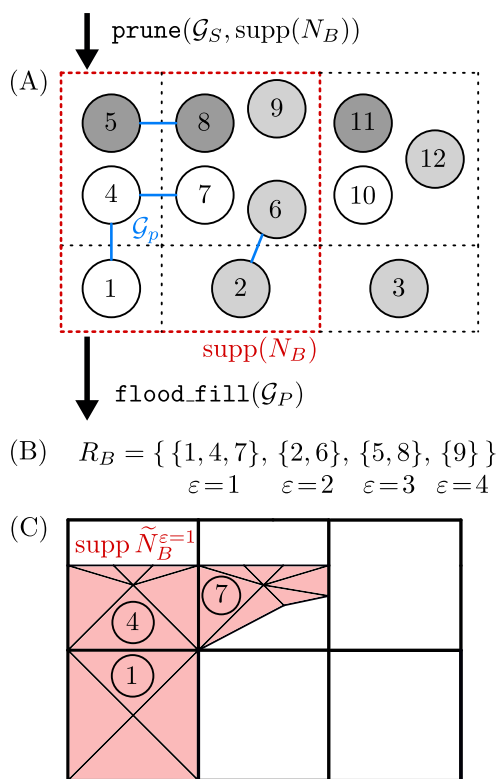


Fig. 20 Basis function enrichment using flood-fill. **A** Restriction of the subphase graph \mathcal{G}_S shown in Figure 19 to the support of a single background basis function N_B . **B** Lists of connected subphases within $\text{supp}(N_B)$ resulting from a flood-fill operation. **C** Support of an example enriched basis function whose support consists of a list of connected subphases

in the cluster data structure introduced in Sect. 3.1. The creation of volumetric *bulk* cluster is discussed first, before the creation of *side* clusters. An example of clusters generated on a single background element is shown in Fig. 22.

For the bulk clusters, the creation is summed up in Algorithm 8. For every subphase in the foreground mesh, the foreground cells are collected. Using the parametric coordinates of their vertices stored on the corresponding child meshes, a standard Gauss-Legendre quadrature rule is mapped into the parametric coordinates of the background element. Pairing the quadrature rule with the unzipped background element associated with the subphase yields a complete cluster. These clusters are grouped by their respective material domains, such that, during assembly, these groups are iterated over, considering the same set of equations associated with the given material domain.

To construct the side clusters and cluster pairs to evaluate integral equations of the form (15) and (16), the steps outlined in Algorithm 9 are necessary. The specific foreground facets F_c between each neighboring but (materially) disconnected subphase pair (S_1, S_2) are determined. Information about the facet connectivity on the foreground mesh \mathcal{F} and the foreground cells c that comprise the subphases S_i are employed.




Fig. 21 Unzipping of background mesh. Copies of background elements are created for each of the subphases contained in them, see Fig. 19. The IEN-arrays of the resulting "unzipped" background elements Ξ_E^u are edited to encode the enrichment information




Fig. 22 Construction of bulk and side clusters in background element $E = 4$ from Figure 19. Left: A connection in the interface subphase graph \mathcal{G}_I corresponds to a pair of side clusters constructed with matching Gauss points on every foreground facet. Right: Mapping Gauss points from the foreground cells to the parametric space of the background elements yields quadrature rules for the volume of every subphase

The ordinals o of these facets with respect to the foreground cells attached to either side are obtained concurrently. Knowing the facet ordinal and foreground cell on either side, Gauss points on the facets are mapped into the coordinate space of the background elements. In addition, the facet has an outward pointing normal, which is stored with the quadrature points.

Algorithm 7 Enrichment: Unzip background elementsInput: Background mesh \mathcal{H} , Foreground mesh \mathcal{T} Output: Unzipped background elements Ξ

```

1: initialize empty container for unzipped background elements  $\Xi \leftarrow \{\}_{E=1}^{n_E}$ 
2: for each background element  $E \in \{1, \dots, n_E\}$  do
3:   size inner container  $\Xi[E] \leftarrow \{\}_{u=1}^{n_u(E)}$ 
4:   for each subphase in the element  $u \in \{1, \dots, n_u(E)\}$  do
5:     copy initialize from background element  $\Xi[E][u] = \text{create\_element}(\mathcal{H}.\text{get\_element}(E))$ 
6:   end for
7: end for
8: initialize enriched basis function counter  $\ell \leftarrow 0$ 
9: for each basis function  $B \in \{1, \dots, n_B\}$  do
10:  initialize list of subphases in  $B$ 's support  $S_B \leftarrow \{\}$ 
11:  for each background element in the support  $E \in \mathcal{H}.\text{get\_elements\_in\_support}(B)$  do
12:    collect subphases in support  $S_B.\text{append}(\mathcal{T}.\text{get\_subphases\_in\_element}(E))$ 
13:  end for
14:   $\mathcal{G}_p \leftarrow \text{prune}(\mathcal{T}.\text{get\_subphase\_graph}(), S_B)$ 
15:  groups of connected subphases  $\{R_B^\varepsilon\}_{\varepsilon=1}^{n_\varepsilon(B)} \leftarrow \text{flood\_fill}(\mathcal{G}_p)$ 
16:  for each enrichment level  $\varepsilon \in \{1, \dots, n_\varepsilon(B)\}$  do
17:    increment the basis function counter  $\ell \leftarrow \ell + 1$ 
18:    for each subphase in the connected group  $S \in R_B^\varepsilon$  do
19:      replace basis function index  $\Xi[S.E][S.u].\text{IEN}.\text{find\_and\_replace}(B, \ell)$ 
20:      (# operation includes safeguard preventing IEN entries from being replaced multiple times)
21:    end for
22:  end for
23: end for
24: assign basis function IDs and get index to ID map  $I_\ell \leftarrow \mathcal{T}.\text{communicate\_basis\_function\_IDs}(\Xi)$ 
25: replace the basis function indices with IDs in IEN arrays  $\text{replace\_IEN\_with\_IDs}(\Xi, I_\ell)$ 
26: return  $\Xi$ 

```

Algorithm 8 Create bulk clusters.Input: Foreground mesh \mathcal{T} , unzipped background elements Ξ Output: Sets of bulk clusters for every material \mathcal{D}_B

```

1: initialize containers of bulk clusters for every material  $\mathcal{D}_B \leftarrow \{\}_{m=0}^{n_m-1}$ 
2: for each subphase  $S \in \mathcal{T}.\mathcal{S}$ , for which  $S.E.\text{is\_owned}()$  do
3:   generate quadrature rule  $\{\xi_q, \omega_q\}_{q=1}^{n_q} \leftarrow \mathcal{T}.\text{map\_param\_gauss\_points}(S.C[:])$ 
4:   create bulk cluster  $\mathcal{D}_B[S.\text{get\_material}()].\text{append}(\text{create\_cluster}(\Xi[S.E][S.u], \{\xi_q, \omega_q\}_{q=1}^{n_q}))$ 
5: end for
6: return  $\mathcal{D}_B$ 

```

With a matched ordering of quadrature points on either side of the interface, the jump terms (8) and (9) are computed without storing additional information. This is achieved by iterating over the attached foreground facets in the same order from either side and then requesting quadrature points from either side in opposite orientation, i.e., clockwise or counter-clockwise. The quadrature rules generated from either side

Algorithm 9 Create side clusters.Input: Foreground mesh \mathcal{T} , unzipped background elements Ξ Output: Sets of side clusters and side cluster pairs for every interface \mathcal{D}_S and \mathcal{D}_I

```

1: initialize containers of side clusters for every material combination  $\mathcal{D}_S, \mathcal{D}_I \leftarrow \{\}_{k=1}^{n_m^2}$ 
2: for each subphase pair  $(S_1, S_2) \in \mathcal{T}.\mathcal{G}_I$ , for which  $S_1.I > S_2.I$  AND  $S_1.E.\text{is\_owned}()$  do
3:   find facets between subphases  $F_S \leftarrow \mathcal{T}.\text{find\_connecting\_facets}(S_1, S_2)$ 
4:   for each side of the interface  $(i, j) \in \{(1, 2), (2, 1)\}$  do
5:     get the cells attached to the facets on this side of the interface with the facet ordinals  $C, o$ 
6:     generated quadrature rule  $\{\xi_q, \omega_q, \mathbf{n}_q\}_{q=1}^{n_q} \leftarrow \mathcal{T}.\text{map\_facet\_gauss\_points}(C[:], o[:], i)$ 
7:     (# the last input indicates whether (counter-)clockwise ordering is used)
8:     ravel the multi-index  $k_i \leftarrow S_i.\text{get\_material}() \cdot n_m + S_j.\text{get\_material}() + 1$ 
9:     create side cluster  $D_i \leftarrow \text{create\_cluster}(\Xi[S_i.E][S_i.u], \{\xi_q, \omega_q, \mathbf{n}_q\}_{q=1}^{n_q})$ 
10:    end for
11:    store side clusters  $\mathcal{D}_S[k_1].\text{append}(D_1), \mathcal{D}_S[k_2].\text{append}(D_2)$ 
12:    store side cluster pairs  $\mathcal{D}_I[k_1].\text{append}((D_1, D_2)), \mathcal{D}_I[k_2].\text{append}((D_2, D_1))$ 
13:    collect non-owned subphases if  $!S_2.E.\text{is\_owned}()$  :  $S_{\text{comm}}.\text{append}(S_2)$ 
14:  end for
15: communicate IENs for non-owned unzipped bg. elements  $\text{communicate\_IEN}(\Xi, S_{\text{comm}})$ 
16: return  $\mathcal{D}_S, \mathcal{D}_I$ 

```

are matched up to the unzipped background elements associated with the subphases on either side to form the side clusters D_1, D_2 and pairs $(D_1, D_2), (D_2, D_1)$, before being stored on the list of side cluster (pairs) associated with the interfaces $\Gamma_{m,l}^I$. The bases \mathcal{B}_E belonging to the different unzipped background elements allow the terms (8) and (9) to be evaluated.

Note on Quadrature Rules

We chose Gauss-Legendre quadrature rules generated on the foreground mesh cells for their simplicity. Other, potentially more efficient, quadrature rules could be generated using the existing subphase information. This could include Stokes' theorem-based methods [34, 35] or moment-fitting [44, 57, 86], both on the volume of the subphases directly or on their boundaries using Stokes' theorem.

3.5 Face-oriented ghost stabilization

This subsection details the implementation of the residual term (19) for face-oriented ghost stabilization. From a theoretical point of view, the challenge here is twofold. The first is to identify the ghost facets \mathcal{F}_G and the adjacent material regions $(u, v) : u \in \{1, \dots, n_u(E^+)\}, v \in \mathcal{U}_{u,E^-}$ for each ghost facet $F \in \mathcal{F}_G$, that is finding the sets of the first three sums in (19). The second is to construct the polynomial exten-

Algorithm 10 Construct side cluster pairs for ghost stabilization

Input: Background mesh \mathcal{H} , Foreground mesh \mathcal{T} , Unzipped background elements Ξ

Output: Ghost side cluster pairs for every material \mathcal{D}_G

```

1: initialize empty container for ghost side cluster pairs  $\mathcal{D}_G \leftarrow \{\}_{m=0}^{n_m-1}$ 
2: for each subphase  $S_1 \in \{1, \dots, \mathcal{T}.G.size()\}$  do
3:   for each neighbor subphase  $S_2 \in \mathcal{T}.G[S_1]$  do
4:     # check if at least one of the adjacent bg. elements is cut and
     # each pair is only constructed once
5:     if ( $\mathcal{T}.n_u[S_1.E] > 1$  OR  $\mathcal{T}.n_u[S_2.E] > 1$ ) AND  $S_1.I > S_2.I$  AND  $S_1.E.is\_owned()$  then
6:       get the connecting facet ordinals  $\{o_1, o_2\} \leftarrow \mathcal{H}.get\_connecting\_facets(S_1.E, S_2.E)$ 
7:       for primary and secondary sides  $i \in \{1, 2\}$  do
8:         get quadrature points on facet  $\{(\xi_q, \omega_q, \mathbf{n}_q)\}_{q=1}^{n_q} \leftarrow get\_quadrature\_points(S_i.E, o_i, i)$ 
9:         (# the last input indicates whether to output clockwise or counter-clockwise ordering)
10:        create side cluster  $D_i \leftarrow create\_cluster(\Xi[S_i.E][S_i.u], \{(\xi_q, \omega_q, \mathbf{n}_q)\}_{q=1}^{n_q})$ 
11:      end for
12:      store cluster pair for material  $\mathcal{D}_G[S_1.get\_material()].append((D_1, D_2))$ 
13:      collect non-owned subphases if  $\neg S_2.E.is\_owned() : S_{comm}.append(S_2)$ 
14:    end if
15:  end for
16: end for
17: communicate IENs for non-owned unzipped bg. elements  $communicate\_IEN(\Xi, S_{comm})$ 
18: return  $\mathcal{D}_G$ 

```

sions $\tilde{\mathbf{v}}^h$ and $\tilde{\mathbf{u}}^h$ and evaluate the jumps in their derivatives at the facets $F \in \mathcal{F}_G$.

The triple sum in Equation (19) is restated as a sum over all subphase neighbors in \mathcal{G}_S , where at least one of the background elements contains two or more subphases, i.e., it contains a material interface. Figure 23 (A) shows an example of this for a single facet $F \in \mathcal{F}_G$ in the mesh in Fig. 19. On this facet, there are four subphase pairs which the second and third sum in (19) amount to. During the unzipping process, outlined in Sect. 3.4, the basis functions on the background

elements \mathcal{B}_E are not changed to include the enrichment function ψ_B^ε . Hence, the basis already represents the polynomial extension $\tilde{\mathbf{v}}^h$ and $\tilde{\mathbf{u}}^h$ without further modification. Pairing up the unzipped background elements corresponding to each of the subphases and adding matching integration points to either side, as is illustrated in Fig. 23 (B), allows for convenient evaluation of the jump terms.

Algorithm 10 outlines the compact procedure to generate the cluster pairs. Analogous to the side cluster construction, a standard Gauss-Legendre quadrature rule is generated on the relevant background element facets with mirrored spatial ordering on either side of each facet. Note, that this is done using the background element’s facet directly, as opposed to collecting the attached foreground cells, reducing the number of quadrature points.

3.6 Parallelization

This subsection addresses the parallel implementation of the presented preprocessing framework. First, the general approach taken to parallelization is discussed before covering the additional considerations for executing each of the previously presented steps in parallel.

The preprocessor uses a domain decomposition approach. The processors independently execute the presented algorithms on parts of the domain assigned to them, allowing for efficient use of distributed memory systems.

An extra layer of elements, referred to as the "aura", is added beyond the part of the background mesh assigned to, or "owned" by, a given processor. This is to keep parallel communication on such systems to a minimum. Due to the enrichment strategy requiring the topological analysis of the full support of each basis function, the aura’s width is such that the support of any basis function, partially supported within a given processor’s owned domain, is fully contained in the owned domain and the aura. The resulting parallel decomposition scheme is illustrated in Fig. 24 using a domain split into four parts and a single basis function. Assuming all basis functions have a support size of 3×3 background




Fig. 23 Ghost side cluster pairs generated from graph. **A** Section of \mathcal{G}_S for a single background facet. **B** Each of the connections in \mathcal{G}_S results in a ghost cluster pair being constructed, since at least one of the adjacent background elements is cut

Algorithm 11 Communicate IDs.Inputs: My proc. p_c , list of adjacent procs. P_{loc}

```

1: # Assign IDs to owned entities
2: initialize O, N = {}
3: for each entity  $e$  do
4:   if is_owned( $e$ ) then
5:     O.append( $e$ )
6:   else # entity is not owned
7:     N.append( $e$ )
8:   end if
9: end for
10:  $\{n_e^i\}_{i=1}^{N_{\text{umProcs}}} \leftarrow \text{MPI\_Allgather}(\text{size}(\text{O}))$ 
11: first ID  $I \leftarrow 1 + \sum_{i=1}^{p_c} n_e^i$ 
12: for each owned entity  $e \in \text{O}$  do
13:    $e.I \leftarrow I + +$ 
14: end for
15:
16: # Send and receive ID requests
17: initialize R, E, A  $\leftarrow \{\{\}\}_{p \in P_{loc}}$ 
18: for each non-owned entity  $e \in \text{N}$  do
19:   find owning proc.  $p_o$ 
20:   get identifying information  $i_d$ 
21:   store request  $R[p_o].\text{append}(i_d)$ 
22:   ... and entity  $E[p_o].\text{append}(e)$ 
23: end for
24: send requests  $\text{MPI\_Isend}(\{R_p\}_{p \in P_{loc}}, P_{loc})$ 
25: receive  $\{R_p\}_{p \in P_{loc}} \leftarrow \text{MPI\_Irecv}(P_{loc})$ 
26:
27: # Answer ID requests
28: for each proc. neighbor  $p \in P_{loc}$  do
29:   for each request  $r \in R_p$  do
30:     find entity  $e \leftarrow \text{find\_entity}(r)$ 
31:     store ID answer  $A[p].\text{append}(e.I)$ 
32:   end for
33: end for
34: send answers  $\text{MPI\_Isend}(\{A_p\}_{p \in P_{loc}}, P_{loc})$ 
35: receive  $\{A_p\}_{p \in P_{loc}} \leftarrow \text{MPI\_Irecv}(P_{loc})$ 
36: for each proc. neighbor  $p \in P_{loc}$  do
37:   for each answer  $i \in \{1, \dots, \text{size}(A_p)\}$  do
38:     store the ID  $E[p][i].I \leftarrow A_p[i]$ 
39:   end for
40: end for

```

elements, like the one shown in the figure, the aura needs to extend by two elements beyond the domain.

An advantage of this approach is that the amount of data communicated between each of the independent processors is minimal. As such, only indexing information within the aura is communicated, as discussed in the next paragraph. While, from a hardware perspective, this allows for efficient scaling across multiple compute nodes, a weakness becomes apparent from the illustration in Fig. 24. If the size of the processors' domains is very small and the aura constitutes a considerable fraction of a mesh handled by a given processor, the computational overhead of processing the aura is large. We are quantifying this overhead in the scaling study presented in Sect. 4.3.

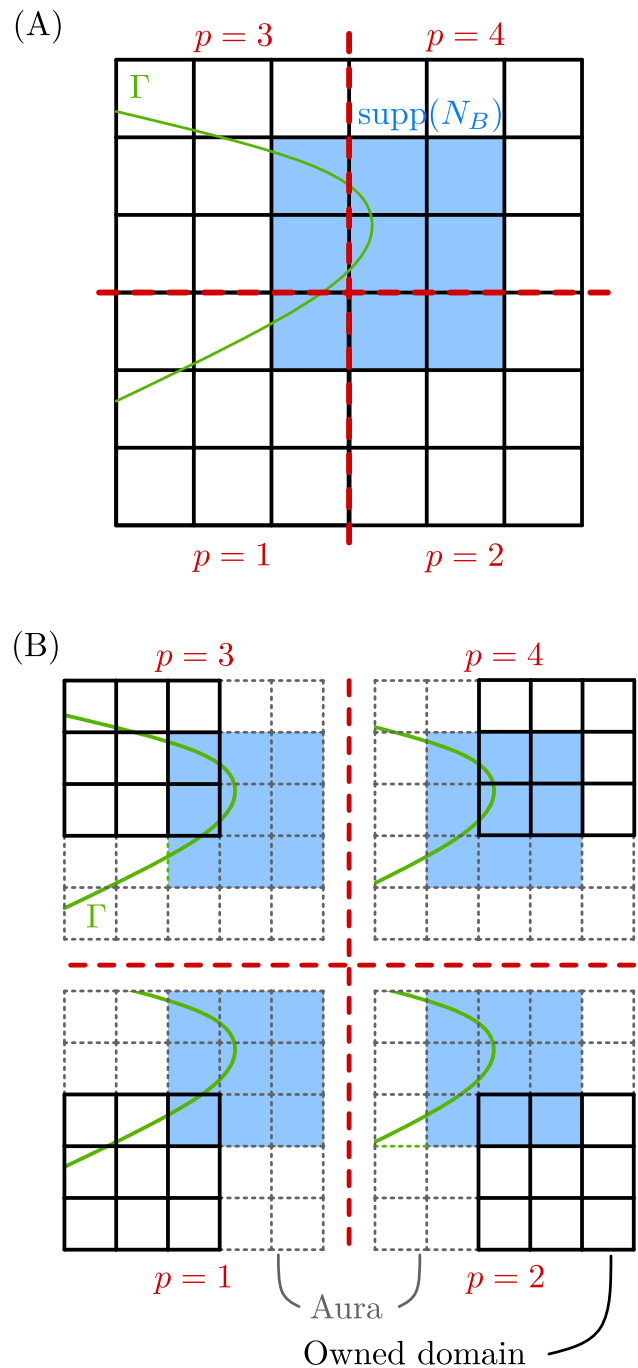
Concept of Parallel IDs

Fig. 24 Domain decomposition in parallel. **A** Global background mesh being split across four processors with the support of one basis function N_B highlighted in blue. **B** Each processor has access to a layer of background elements, the "aura", beyond the owned domain such that the support of any basis function in the owned domain can be analyzed for enrichment

Entities, that are present on two or more processor domains simultaneously, need to be handled in a parallel consistent manner. As the indices of entities differ between processors, a parallel consistent, and globally unique ID is assigned to each entity. Further, it is assumed that any entity is owned

by exactly one processor. Upon creation of entities of any type, Algorithm 11 is used to assign IDs to the owned entities and communicate IDs with neighboring processors that have access to those entities but do not own them. The first step of the procedure, i.e., the assignment of IDs to owned entities, only requires the communication of the number of owned entities by each processor. In the second step, information is collected for non-owned entities with which they are identified unambiguously by the owning processor. The content of this information is discussed in the following. In the last step, using the identifying information received from adjacent processor domains, the IDs of the associated entities are collected and sent back to the requesting processors. The communication calls in Algorithm 11 are adapted from the OpenMPI library [32].

Background Mesh

For the presented algorithms to work in parallel, the data structure of the input mesh, shown in Fig. 8, needs to meet the following requirements:

- The list of elements \mathcal{E} provided to a given processor p includes both the owned and the aura elements. Each element $E \in \mathcal{E}$ stores its owning processor p and has a parallel ID assigned to it.
- The basis functions provided by the mesh and their indexing are restricted to those basis functions that are at least partially supported within the owned domain. Basis functions only supported within the aura are ignored. The IEN arrays of the aura elements are left empty.
- All basis functions are assigned an owner and a parallel consistent ID.

Foreground Mesh

The foreground cells and the subphases need parallel consistent IDs. The parallel consistent subdivision in each processor's aura is ensured by the order in which the `for` loops in Algorithm 3 are executed. After the regular subdivision templates are applied in identical orientation by each processor, the foreground cells generated during the subdivision are stored in identical order and with identical orientation on the child meshes across the processors. Knowing the ID of the background element a child mesh is located in and the local index of a given cell inside the child mesh, cells are identified across processors and their IDs communicated in the function `T.communicate_cell_IDs()` in Algorithm 1. Subsequently, subphases are identified by providing the ID of one of the foreground cells as part of the function `T.communicate_subphase_IDs()` in Algorithm 5.

Enrichment

The enriched basis functions require parallel consistent IDs to assemble the global system of equations. As part of the

enrichment process in Algorithm 7, one of the subphases comprising the support of an enriched basis function is collected to identify enriched basis functions across processors. The basis function indices listed in the IEN arrays of the unzipped background elements are replaced with their respective parallel IDs.

Clusters

Generally, only clusters and cluster pairs contained within a processor's domain are constructed to ensure each cluster or cluster pair's contribution is only added once to the global linear system during assembly. For side cluster pairs, particularly for those constructed for the purpose of ghost stabilization, situations may arise, where the two adjacent unzipped background elements are part of different processors' domains. In this case, the adjacent subphases are collected in Algorithms 9 and 10, and the content of the IENs of the associated unzipped background elements is communicated. Recall, that the IEN of background elements within the aura was initially left empty.

4 Numerical examples

This section presents a series of numerical examples which study the features and performance of the preprocessing framework and its implementation specifically. For a detailed analysis of the numerical framework's performance, e.g., its convergence behavior and stability properties, the reader is referred to [62]. The first example in Sect. 4.1 is used to validate and discuss the impacts of the enrichment strategy and the ghost stabilization for a geometry with small features. The "Brick Wall" example in Sect. 4.2 demonstrates the preprocessing framework's robustness for pathological geometric cases with multiple geometries, sharp features, vanishing material domains, and coinciding interfaces. Lastly, we demonstrate the performance and parallel scalability with a "scaffold sandwich" structure in Sect. 4.3.

While the numerical examples presented in the following are all structural problems governed by the equations introduced in Sect. 2.1, the preprocessing framework applies to most types of (continuum) physics – including those with non-linear governing equations – that do not involve interfaces that evolve in time. Earlier works have applied the implementation presented herein to, e.g., thermo-structural problems [62, 75] or to incompressible fluid flow with thermal advection–diffusion [61].

The condition numbers presented were estimated using either Matlab's `cond()` function for small systems or by solving for the smallest and largest eigenvalues using a Krylov-Schur Method [85], as implemented in the SLEPC package [39]. To solve linear systems, the direct solver implemented in PARDISO [71, 72] and a flexible generalized




Fig. 25 A Setup of the example with multiple connected beams including the background mesh and the support of a single basis function. **B** Normal stresses in the solution to the non-enriched problem and **C**

the problem using the enrichment strategy (11) with fillets added to the interior corners. The reference displacement solution generated using a highly refined mesh is shown in gray in the background

minimal residual solver (FGMRES) [68] as implemented in the PETSc library [9] were used.

4.1 Multiple thin beams

We consider the setup shown in Fig. 25 (A) with multiple connected beams under a bending load. The beams are immersed in a coarse background mesh with element size $h_0 = 1$. The material has a Poisson ratio of $\nu = 0.3$ and Young’s modulus of $E = 1000$. A uniform body load $\mathbf{b} = (1, 0)$ is applied in x -direction. For the initial setup, the beams have a width of $\delta_t = h_0$, are spaced apart by $\delta_d = h_0$, and, to avoid generating a body-fitted mesh, are offset by $\delta_o = 0.6 h_0$ relative to the background mesh. The problem is discretized using quadratic C^1 -continuous B-spline basis functions.

In this configuration, without basis function enrichment,⁴ the support of individual basis functions spans across multiple features, as shown in Fig. 25 (A). This results in severe artificial stiffening, or "cross-talk", as the basis functions incorrectly couple the individual beams. This results in the normal stress, σ_{yy} , being under-predicted by a factor of approximately 5.5, see Fig. 25 (B). The enrichment strategy (11) remedies this issue by effectively creating separate enriched basis functions for each beam. Even with the excessively coarse background mesh, the enriched solution follows the overall stress response of the reference solution created using a highly refined mesh, as shown in Fig. 25 (C).

The impacts of the enrichment strategy can further be quantified by comparing the error in the solution with and without enrichment under quad-tree, i.e., isotropic, background mesh refinement. Omitting enrichment, the errors in the stresses, see Fig. 27, drastically decrease as the size of




Fig. 26 Relative L2-error with increasing mesh resolution with and without the generalized enrichment strategy using quadratic $p = 2$ B-splines. Note that the convergence rate r is double the slope when plotted against the number of DOFs in 2D

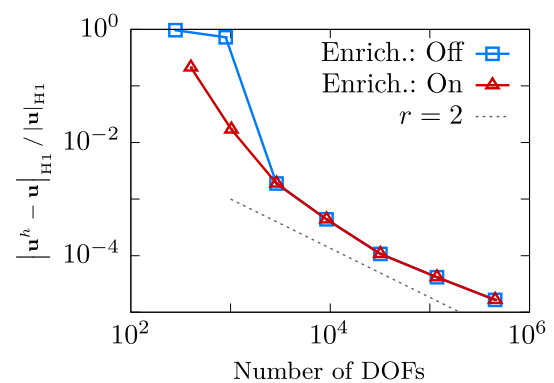


Fig. 27 Relative H1-semi-error with increasing mesh resolution with and without the generalized enrichment strategy using quadratic $p = 2$ B-splines

⁴ The numerical framework without basis function enrichment is equivalent to CutFEM [16] with B-spline basis functions, sometimes also referred to as CutIGA [25].

the support of basis functions shrinks to below the distance between the beams δ_d ; only once the mesh is sufficiently

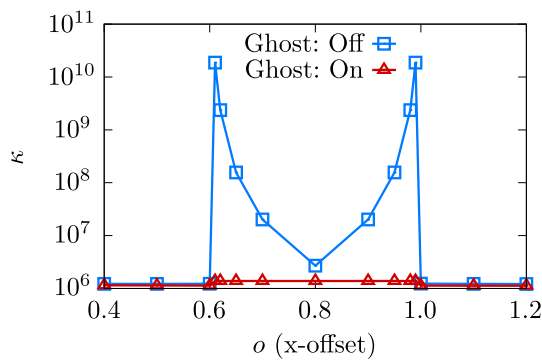


Fig. 28 Condition number κ of the linear system with and without the face-oriented ghost stabilization applied as the structure is translated in x -direction through the background mesh. A penalty of $\gamma_G^k = 0.01 \cdot Eh_0^{2k-1}$ is used in the stabilized case




Fig. 29 Geometry of the "Brick Wall" example

refined, the errors converge at an (above) optimal rate.⁵ Employing the proposed enrichment strategy, we do not observe the initial region of high error and below optimal convergence due to under-resolved features. Note that the enriched problem has a greater number of degrees of freedom compared to the non-enriched problem on an equal background mesh with element size h . For higher mesh resolutions, the enrichment has no impact on the number of degrees of freedom and the error in the solution. The enrichment, in essence, adds anisotropic refinement in regions where, and only where, the effects of the material topology would otherwise not be captured by the background mesh. Hence, the generalized Heaviside enrichment strategy loosens the requirements on the mesh resolution relative to the feature size in a given problem. Unlike a local refinement-based approach to mitigate cross-talk [20], the enrichment scheme adds fewer additional degrees of freedom and does (most likely) not affect the critical time step size in, e.g., fluid dynamics problems. From a conceptual standpoint, the enrichment strategy and the control point duplication approach presented in [49] are identical; hence, they exhibit the same overall behavior in terms of where degrees of freedom are added and their effect on the critical time step size.

Ghost Stabilization

We investigate the impact of the ghost stabilization on the enriched problem by considering the setup shown in Fig. 25 (A). Compared to the previous study, the beam spacing is decreased to a size smaller than the background element size $\delta_d = 0.6 h_0$, and the beam width is increased to $\delta_t = 1.4 h_0$, such that multiple beams connect in the same background element. The configuration requires the construction of multiple ghost penalty terms on the same facet, as shown in Fig.

6. The position of the beams δ_o relative to the background mesh is then shifted to obtain various cut configurations. The impact on the condition number is shown in Fig. 28. The condition number for the non-stabilized system spikes as the beams' edges transition through the background mesh lines, while the stabilized system only experiences small changes in the condition number due to the double penalization of a single facet leading to an increased contribution of the ghost penalty term. The conditioning of the linear system remains unaffected by the location of intersections besides this small perturbation. This aligns with the expected behavior for ghost stabilization and indicates that the adapted ghost stabilization scheme successfully targets any insufficiently supported enriched basis functions.

4.2 Brick wall

This example aims to demonstrate the robustness of the pre-processing framework in a multi-material setting, specifically in terms of handling geometric edge cases and in terms of retaining good conditioning of the linear system. To this end, a multi-material geometry involving straight edges, vanishing material domains, and coinciding interfaces is considered and randomly placed within a background mesh.

The geometric setup consists of a $2 \times 1 \times 1$ cuboid cut in half in each spatial dimension by three axis-aligned planes. Additionally, two planes are placed symmetrically at varying positions $x = -\delta$ and $x = +\delta$ from the center of the cuboid. Some of the resulting subdomains are then joined together to generate material blocks, or "bricks", of various shapes, as shown in Fig. 29. Examples of various edge cases in the alignment of the geometry with the background mesh and the resulting foreground meshes generated by the preprocessor are shown in Fig. 30.

The impact of the geometry's position in the background mesh on the conditioning of the linear system is determined

⁵ To mitigate stress singularities and obtain optimal convergence rates, fillets are added to the interior corners, as shown in Fig. 25 (B),(C). Additionally, the foreground mesh is refined to limit the geometric error.

by immersing the geometry in a Cartesian grid background mesh with element size $h = 0.1$, such that for $\delta = 5h$ all interfaces align with the background mesh lines. From this starting configuration, the distance δ is reduced to 0 to generate increasingly smaller material domains. In addition, the geometry is translated randomly in all coordinate directions 100 times for each value of δ , where the translation in each spatial dimension is given by

$$\Delta x_i = \text{sign}(Z) \cdot 10^{-|Z|} \cdot h, \quad (24)$$

with Z being a random number drawn from a normal distribution with mean 0 and standard deviation 1. Quadratic C^1 -continuous B-spline basis functions are used for discretization, and random, uniformly distributed, Young's moduli $E \in [1, 5]$ and Poisson ratios $\mu \in [0.25, 0.35]$ are assigned to the different material blocks. The whole structure is clamped at $x = -1$ and a traction load of $\mathbf{t} = (1, 1, 1)$ is applied at $x = 1$. The condition number ranges of the linear system resulting from the different cut configurations, with and without ghost stabilization, are shown in Fig. 31.

The results demonstrate that the problem exhibits a conditioning behavior expected for the face-oriented ghost stabilization: the stabilization method keeps the conditioning close to the best-case cut configuration for the given geometry, independent of the geometry's position in the background mesh. For small values of δ , thin slices of material are created which are only supported by cut background elements. Hence, the ghost stabilization is not able to stabilize cut basis functions against non-cut basis functions and prevent a slight increase in the condition number κ . The system remains solvable for the considered material sliver sizes with the linear solvers employed in this study. However, this may not remain true for smaller sliver sizes. This observation informs the choice of numerical tolerances that should be used in the `compute_proximity` and `find_interface` functions in Algorithm 3. Tolerances for the proximity of a vertex to be considered "on the interface" ($P(v) = 0$) should be chosen large enough to ensure numerical stability in geometric edge cases, like the one presented, as opposed to only being large enough to avoid errors due to machine precision arithmetic.




Fig. 30 Edge cases for cut configurations and resulting foreground mesh on a $2 \times 2 \times 2$ element section of the bg. mesh, shown in black. **A** Intersections align with the centers of bg. elements. **B** Intersections align with the edges of bg. elements. **C** Rotation of the geometry in the bg. mesh with interfaces coinciding at a background vertex. **D** Small offset of interfaces to the bg. mesh resulting in slim material slivers

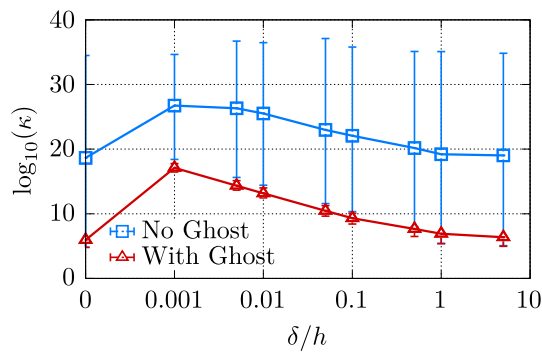


Fig. 31 Mean and range of $\log_{10}(\kappa)$ of the linear systems resulting from the "Brick Wall" example for varying cut distances r and random translations of the geometry inside the background mesh. The stabilization penalty is set to $\gamma_G^k = 0.01 \cdot Eh^{2k-1}$. Note, the left-most entry violates the scale and is $\delta/h = 0$

4.3 Scaffold sandwich

We consider the scaffold structure with thin geometric features, shown in Fig. 32 (A), to investigate the computational performance and scalability of the preprocessing framework. The structure consists of a periodic grid of unit blocks shown in subfigure (B). This problem setup is chosen as the size of the numerical model can be altered by increasing the number of unit blocks, whereas the global geometric complexity – the portion of background elements intersected by an interface, the number of enriched basis functions generated relative to the number of basis functions on the background mesh, the density of the skeleton mesh of ghost facets, etc. – stays constant. For the following parallel scaling studies, the domain is only decomposed in x - and y -directions to ensure that the geometry near processor interfaces remains similar. Further, the effect of the parallel domain decomposition on the scaling behavior due to unequally distributed geometric complexity is mitigated by ensuring that each processor’s domain contains approximately the same geometry.

Each unit block is discretized using a part of the background mesh with $24 \times 24 \times 96$ elements and quadratic C^1 -continuous B-spline basis functions. Each unit block contains ~ 7400 intersected background elements and $\sim 530,000$ cells in the foreground mesh.⁶ A detailed view of the tessellation generated is shown in Fig. 32 (C).

Three types of scaling are investigated.

1. **Time complexity:** The problem size is increased by increasing the number of unit blocks while keeping the number of processors constant.
2. **Weak parallel scaling:** The problem size and the number of processors are increased simultaneously. The weak

⁶ A voxel-based geometry definition for the entire structure is used which does not align exactly with the background mesh; hence, the geometry and number of cut elements in each unit block may differ slightly.

scalability is an indicator of how well the framework performs in scenarios where both the computational resources available and the problem size are large.

3. **Strong parallel scaling:** The number of processors is increased while keeping the problem size constant. Strong scalability is a crucial factor in how viable the framework is in use cases where the solution time is critical, and a large number of processors is available. Examples of such scenarios are optimization or design studies during which many designs are analyzed in rapid succession.

Quantifying Scalability

The generalized efficiency measure η is used to quantify the scalability, i.e., how efficiently the framework uses computational resources in the three scaling scenarios.

$$\eta = \frac{\frac{s}{s_0}}{\frac{n_p}{n_{p,0}} \cdot \frac{t}{t_0}}, \tag{25}$$

where s , n_p , and t denote the problem size, the number of processors used, and the time to perform the various preprocessing steps, respectively. The subscript $(\cdot)_0$ denotes the baseline values relative to which the efficiency is evaluated. An efficiency value of $\eta = 1$ indicates perfect scaling. The measure is generalized in the sense that it is applicable to evaluate strong and weak parallel scalability, as well as problem size scaling assuming that the algorithms in question scale linear in time, $t = \mathcal{O}(s)$. Additionally, the memory efficiency is defined as

$$\mu = \frac{\frac{s}{s_0}}{\frac{M}{M_0}}, \tag{26}$$

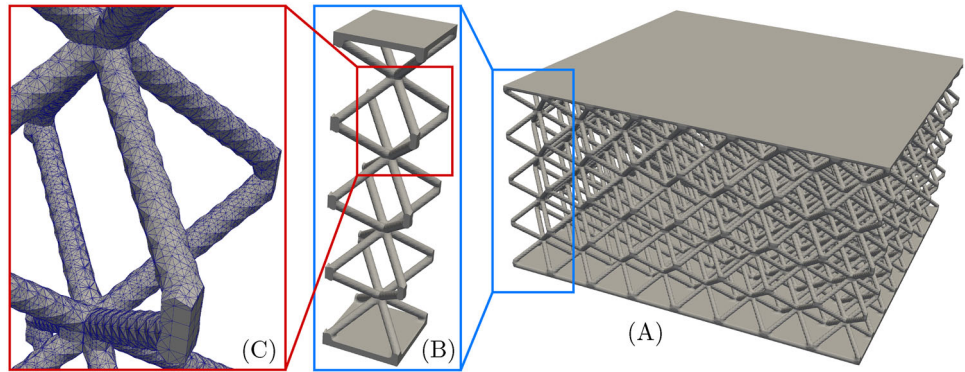
where M denotes the memory consumption.

We define two mesh overhead values to better understand the scaling results and how they are impacted by the chosen parallel strategy. The processor local mesh overhead λ_{loc} is defined as the maximum ratio across all processors of background elements on a processor including the aura over the number of elements owned by it

$$\lambda_{loc} = \max_{p \in P_{glob}} \left(\frac{\text{\#bg. elems. on } p}{\text{\#bg. elems. owned by } p} \right). \tag{27}$$

Meanwhile, the global mesh overhead λ_{glob} is the total number of background elements processed across all processors – counting aura elements multiple times as they are shared between different processors – over the actual number of

Fig. 32 Geometry for the scaffold sandwich. **A** Largest problem size constructed from an 8×8 grid of unit blocks. **B** Single unit block which is periodically repeated to construct the example. **C** Detail of the tessellation generated during the decomposition



background elements.

$$\lambda_{glob} = \frac{\sum_{p \in P_{glob}} \#bg. \text{ elems. on } p}{\sum_{p \in P_{glob}} \#bg. \text{ elems. owned by } p}. \quad (28)$$

The efficiency measures η and μ are compared to the local or global mesh overheads relative to the baseline configurations $\lambda' = \lambda/\lambda_0$.

Hardware and Measurements

All results were run on a single server node using a single AMD EPYC™ 7713 64-core processor with an eight-channel DDR4 memory configuration. Dynamic frequency scaling was disabled to prevent any parallel efficiency losses due to clock-speed variations.

The timing and memory consumption of the presented preprocessor is split into three sections: the *decomposition*, *enrichment*, and *ghost facet generation* consisting of the steps described in Sect. 3.2 and Sects. 3.3, 3.4, and Sect. 3.5, respectively.

The memory consumption is measured using the memory allocation reported by the operating system. This measurement slightly overestimates the actual amount of memory used, though this error is negligible when measuring memory consumption in the decomposition and enrichment steps. However, the creation of the ghost stabilization facets consumes comparatively negligible amounts of memory, on the order of $1e0 \sim 1e1$ megabytes for problems consuming $1e1 \sim 1e2$ gigabytes, which cannot be measured accurately. Given their insignificance, these measurements are omitted from the results. Running the large-scale problems presented with a memory profiler to obtain precise memory consumption would be cost-prohibitive. For runtime measurements, the maximum core time across all processes is used.

All measurements are averaged over three independent runs of each problem setup. The baseline values for timing and memory consumption, t_0 and M_0 , for the results presented in the following, are summarized in Table 1.

Time & Memory Complexities




Fig. 33 Runtime efficiency η when scaling up the problem size from one to 64 unit blocks using four processors

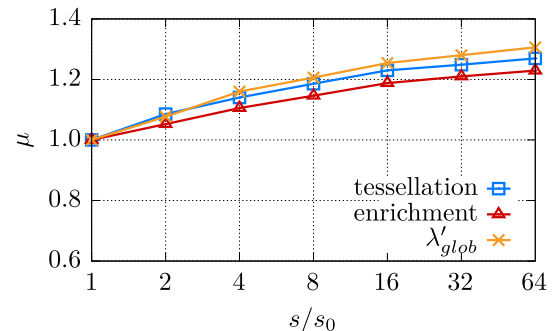


Fig. 34 Memory efficiency μ when scaling up the problem size from one to 64 unit blocks using four processors

The time and memory complexities of the framework are analyzed by scaling the problem size from a single unit block to a grid of $8 \times 8 = 64$ unit blocks. The problems were run in parallel on four processors with the domain split in half in x - and y -directions in each case. The runtime and memory efficiencies η and μ are shown in Figs. 33 and 34, respectively. The relative mesh overhead decreases as the problem size increases, given that the number of processors is constant. Hence, efficiency measures $\eta, \mu > 1$ are observed.

The runtime for the decomposition closely follows λ'_{loc} which indicates that the algorithms employed have an approximately linear time complexity $t \approx \mathcal{O}(s)$. For the ghost facet generation, $\eta \approx 1$ throughout. Given that the individual processors do not generate ghost facets in the

aura, these operations are not expected to be affected by the decreasing relative mesh overhead. For the enrichment, only part of the basis functions supported in the aura is processed. As the runtime for the enrichment scales between the mesh overhead and 1.0, the operations have an approximately linear time complexity $t \approx \mathcal{O}(s)$.

The memory consumption of the decomposition step closely follows the expected memory consumption due to the mesh overhead. The enrichment is less affected by the mesh overhead, and therefore its memory efficiency remains slightly below λ'_{glob} . These results suggest that the memory consumption of the algorithms scales linearly with the problem size $M = \mathcal{O}(s)$.

Weak Parallel Scaling

The weak parallel scalability is investigated by increasing both the number of parallel processors and the number of unit blocks in the problem from one to 64, yielding a domain size of one unit block per processor.

The weak parallel efficiency is shown in Fig. 35. For the decomposition, η mostly follows the relative mesh overhead with a few percentage points offset, which can be explained by the approximately constant time it requires to perform the communication procedure in Algorithm 11. The mesh overhead increases as new processor boundaries and, therefore, aura elements are created when adding unit blocks to the initial 1-processor case with no aura. Notably, though, the runtime efficiency decreases significantly when scaling to above 16 processors. This is caused by a saturation of the available memory bandwidth. This assumption was confirmed by repeating the study on a server node with half the number of memory channels, and therefore memory bandwidth, but otherwise identical hardware. Its parallel scaling behavior showed the same significant performance breakdown scaling above 8 processors. The enrichment is not as strictly affected by the mesh overhead but suffers from the same problem. It is a known problem that the improvements to raw compute performance have outpaced improvements to memory performance for multiple decades of high-performance hardware development [53, 80]. The result is that for modern hardware, such as the one used here, the memory performance is generally the bottleneck. Additionally, the presented algorithms heavily rely on storing and accessing data at different points in the workflow, which results in cache misses and frequent random memory access operations. Hence, the algorithms presented end up being memory-bound.

Memory consumption can, again, accurately be predicted by the mesh overhead as shown in Fig. 36.

Strong parallel scaling

We investigate the strong Parallel Scalability by increasing the number of processors from one to 64 on two problems, a larger one with 64 unit blocks and a smaller one with 16 unit




Fig. 35 Weak scaling efficiency η when increasing the number of both the processors and unit blocks in the problem from one to 64




Fig. 36 Memory efficiency μ when increasing the number of both the processors and unit blocks in the problem from one to 64

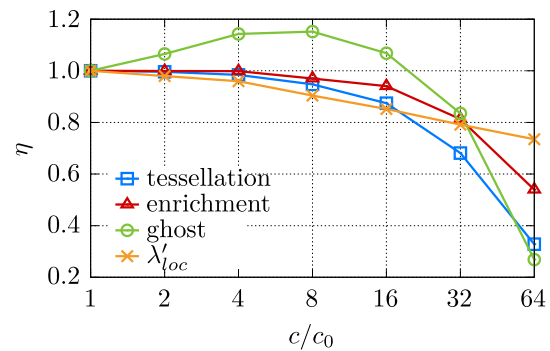


Fig. 37 Strong scaling efficiency η when increasing the number of processors from one to 64 on a **large** problem consisting of 64 unit blocks




Fig. 38 Strong scaling efficiency η when increasing the number of processors from one to 64 on a **small** problem consisting of $4 \times 4 = 16$ unit blocks

Table 1 Runtime t_0 and memory consumption M_0 for the baseline cases in Figs. 34–38

scaling type	problem size	weak	strong (large)	strong (small)
tessellation t_0	1.740s	4.801s	349.5s	84.05s
enrichment t_0	1.895s	6.097s	439.8s	110.6s
ghost facet gen. t_0	0.084s	0.265s	24.71s	5.41s
$\lambda_{loc,0}$	1.361	1.0	1.0	1.0
tessellation M_0	1250 MB	919 MB	60775 MB	13311 MB
enrichment M_0	1621 MB	1239 MB	81519 MB	29395 MB
$\lambda_{glob,0}$	1.333	1.0	1.0	1.0

blocks. Timing results for these two cases are shown in Fig. 37 and 38, respectively. Scaling to more than 16 processors, again, results in a significant performance breakdown due to memory bandwidth saturation, regardless of the problem size. Ignoring this, the runtime efficiency for the decomposition follows the mesh overhead, while the enrichment is not as affected by the mesh overhead. Compared to the weak scaling scenario, the overhead created by communication is negligible. Comparing the two size cases, the decreasing performance ceiling due to the mesh aura for smaller problem sizes presents itself. E.g., for the smaller problem, the overall parallel efficiency for running the smaller scale problem on 16 processors is limited to only $\eta \approx 80\%$, as opposed to $\eta \approx 91\%$ for the larger problem.

Summary

Both the compute time and memory consumption for the presented preprocessing framework scale approximately linearly with the problem size if the geometric complexity is kept constant. In weak and strong scaling scenarios, the compute time and memory consumption scale close to optimally, once the mesh overhead is accounted for, with only a couple of percentage points decrease due to communication in the weak scaling scenarios.

The memory-bound design of the algorithms and the mesh overhead created by the aura present arguably the biggest shortfalls of the preprocessing framework. Re-designing the algorithms to more smartly (re-)compute data in place of access would mitigate the former bottleneck. However, it remains an open question to what extent this is possible.

The mesh overhead can constrain the performance severely when moving to very small background meshes on each processor's domain. As an example, consider a problem using quadratic C^1 -continuous B-splines and running on a large number of processors such that the processors' domains are surrounded by those of adjacent processors to all sides. Reducing the processor-local mesh size to, e.g., $10 \times 10 \times 10$ elements results in a mesh overhead limiting the performance to $\eta \approx \lambda'_{loc} = 2.6\%$, whereas a local mesh size of $25 \times 25 \times 25$ elements leads to a more modest loss $\eta \approx \lambda'_{loc} = 64\%$. Given the overall speed of the procedures presented, it is likely unnecessary to move to

such small processor-local domains during preprocessing. Re-distributing the generated clusters among a larger number of processors after preprocessing would be a solution if small processor-local domains are necessary to rapidly perform element assembly.

5 Conclusion

In this article, we discussed the implementation of an efficient, parallel, and robust immersed finite element preprocessing framework, including the underlying algorithms and data structures.

A tessellation procedure creates a grid of cells fitted to the interfaces and generates quadrature rules on the different material regions and their interfaces. Material connectivity data is built on this grid and provides the information necessary to subsequently perform basis function enrichment and apply a stabilization scheme. This approach of building information about the material connectivity from a geometric decomposition of the background mesh forms the backbone of the pre-processing framework. It allows some of the challenges associated with immersed finite element methods – the generation of custom quadrature rules, enrichment if needed, and stabilization – to be addressed cohesively.

The tessellation procedure applies subdivision templates repeatedly for an arbitrary number of geometries, enabling multi-material domains to be defined. Further, the presented algorithms for building topological information are designed without floating-point arithmetic, making them highly robust. The preprocessor's output is organized into elements paired with custom quadrature rules. Any information about the enriched finite element basis and stabilization methods is implicitly contained in them and not further exposed to the outside as such. Hence, assembly may be performed with existing, standard finite element routines without or with only minimal modifications.

The presented examples demonstrate the robustness and performance characteristics of the preprocessor. The generalized enrichment strategy prevents spurious coupling between close but disconnected features and thereby eliminates the case where the lack of mesh resolution leads not only to a

large error but also to the simulation of a geometrically different, de-featured domain. Robustness is demonstrated in multi-material cases, both in terms of handling geometric edge cases and in terms of retaining good conditioning. The implementation's time complexity and memory consumption are found to scale approximately linearly with the problem size η , $\mu \approx \mathcal{O}(s)$, and the overhead created by parallel communication is shown to be minimal.

However, as summarized at the end of Section 4.3, the outlined parallel strategy creates a performance overhead for very small processor domains. Further, the scaling study highlighted that the implementation is limited by memory performance, rather than compute performance due to the algorithms' reliance on frequently accessing data stored in memory. Future work substituting the tessellation procedure and quadrature method, which only yield a linear geometry approximation, with ones that capture the geometry with higher-order accuracy, would likely improve the framework's numerical efficiency. This especially applies to problems involving non-linear interface phenomena, such as those encountered in contact mechanics, where highly refined meshes are currently needed. Furthermore, the framework, in its current form, does not support interfaces and external boundaries evolving in time. While mesh transformation-based approaches have been applied to enable, e.g., the simulation of fluid–structure problems with large deformations [43], due to the required mesh transformation, they contend with mesh quality issues in the case of (very) large deformations and an inability to handle topological changes [46]. The arguably most natural extension to the presented framework, enabling the simulation of problems with dynamically evolving geometries, would be to also discretize the geometry in time, yielding a space-time approach similar to those presented in [38, 58].

Appendix A Glossary

Adjacency list • Data structure to represent the graphs $\mathcal{G}_{(\cdot)}$ as an array of arrays. In our case, the outer array entries correspond to each subphase and the inner (unordered) array contains the adjacent subphase neighbors, in the sense of the particular graph.

(Background) ancestry/ancestor • Symbol: a • The background ancestor of an entity of the foreground mesh is the background entity of the lowest rank, i.e., dimensionality, which fully contains the foreground entity. See: Section 3.2, par. Background Ancestry.

Background mesh • Symbol: \mathcal{H} • Input mesh constructed on a geometrically simple domain into which the geometry to be analyzed is immersed. It defines the (non-enriched) finite element basis. The associated data structure is shown in Fig. 8.

Cell • Symbol: c • Purely geometric, d -dimensional entity on a d -dimensional mesh. They are, e.g., the triangles and quadrilateral in 2D, or tetrahedrons and hexahedrons in 3D comprising the foreground and background meshes. Note, that unlike elements, cells do not have a basis associated with them. See: Section 3.1, par. Mesh Entities and Connectivity.

Child mesh • Symbol: CM • Group of foreground cells generated from a single background element during the tessellation procedure outlined in Section 3.2. The associated data structure is shown in Fig. 13.

Cluster • Symbol: D • Data structure that combines a (background) element with a (custom) quadrature rule defined on its domain, or a subdomain thereof. The output of the preprocessors are groups of clusters, intended for standard finite element assembly routines. The associated data structure is shown in Fig. 8. The concept of a cluster is discussed from a theoretical standpoint in Section 2.2.

(Parallel, mesh) decomposition • Process of dividing the background mesh into subdomains to be handled by processors independently. See: Section 3.6.

(Background, finite) element • Symbol: E • A finite element consists of a domain Ω^E and a set of basis functions \mathcal{B}_E with their associated degrees of freedom IEN. In the article, the only finite elements used are those comprising the background mesh. Hence, they are referred to as background elements. The associated data structure is shown in Fig. 8.

Enrichment function • Symbol: ψ • The indicator function associated with an enriched basis function with multi-index $\ell = (B, \varepsilon)$. For a definition, see (12).

(Mesh) entity • Symbol: e • Geometric objects comprising a mesh or grid, namely: vertices, edges, faces, 2D/3D cells. See: Section 3.1, par. Mesh Entities and Connectivity.

Facet • Symbol: F • The $(d - 1)$ -dimensional entity of a d -dimensional mesh, e.g., edges in 2D and faces in 3D. See: Section 3.1, par. Mesh Entities and Connectivity.

Foreground Mesh • Symbol: \mathcal{T} • Body-fitted mesh resulting from the subdivision process outlined in Section 3.2. Additionally, the associated data structure, shown in Fig. 13, carries information about the material topology constructed in Section 3.3.

Geometry • Symbol: G • Object defining a material interface or volumetric domain. The associated data structure is shown in Fig. 8.

(Parallel) ID • Symbol: I • An entity ID is a parallel-globally unique identifier for a mesh entity that ensures consistency and facilitates communication across processors in parallel computations. See: Section 3.6, par. Concept of Parallel IDs.

Index • A (processor-local) identifier for entities or other programming objects of a specific type, ranging from 1 to the total number of such entities or objects on a processor. Indices are used to access entities or objects that are stored in arrays. See: Section 3.1, par. Mesh Entities and Connectivity.

Initializer list • A list of values, entities, and/or objects that are used to create/initialize another entity or object.

Ordinal • Symbol: o • The local index of an entity with respect to another entity, as shown in Fig. 10 (A). See: Section 3.1, par. Mesh Entities and Connectivity.

Proximity • Symbol: P • Whether an entity lies outside (+), inside (−), or on the boundary (0) of the geometry within some tolerance. The concept is first introduced in Section 3.1 and used extensively in the step of material assignment.

Queue • Symbol: Q • Data structure collecting new entities/objects to be created, also ensuring their uniqueness. The concept is discussed in Section 3.2.1. The associated data structure is shown in Fig. 40.

Rank • Symbol: r • Dimensionality of an entity in the foreground mesh. See: Section 3.1, par. Mesh Entities and Connectivity.

Subphase • Symbol: S • Topologically connected material region within a single background element’s domain Ω^E . Their generation is discussed in Section 3.3, par. Identification and Generation of Subphases. data structure: Fig. 17.

Tessellation • Process of dividing a domain into geometric primitives. In the context of this article, it is used to refer to the subdivision process generating the foreground mesh.

Unzipping • Index: u • Process of duplicating background elements intersected by material interfaces and re-indexing their basis functions to represent the enriched basis. The process is described in Section 3.4, par. Unzipping and illustrated in Fig. 21.

Vertex • Symbol: v • Purely geometric 0-dimensional entity on a mesh. They are the corners (or endpoints) of cells (or edges). Note, that unlike nodes, vertices do not have the notion of a basis function associated with them. See: Section 3.1, par. Mesh Entities and Connectivity.

Appendix B Additional Algorithms and Data Structures

B.1 Determining Ancestry of New Vertices

Using the particular regular subdivision templates shown in Fig. 14, the background ancestry of a new vertex added at any stage during the hierarchical templated subdivision process outlined in Section 3.2.2 can be determined by the ancestry of the two vertices attached to that edge. One only needs to check which of the following conditions is true:

1. One of the vertices is inside the bg. element: the new vertex is inside the element.
2. The edge vertices descend from the same edge or face: the new vertex is a descendant of that edge or face.
3. Both edge vertices each descend from a background vertex: these bg. vertices are necessarily part of the same

Algorithm 12 Generate entity connectivity.

Input: mesh \mathcal{T} containing a cell-vertex connectivity \mathcal{C} , rank of entities r

Output: entity connectivity \mathcal{F}

```

1: initialize empty entity connectivity  $\mathcal{F}$  and hash map  $VtE : (v_1, v_2, \dots) \mapsto e$ , and running index  $i \leftarrow 0$ 
2: # the number of cells is known and the cell-to-entity array in  $\mathcal{F}$  can be initialized accordingly
3: for each cell  $c \in \mathcal{T.C}$  do
4:   for each entity ordinal on cell  $o \in \{1, \dots, c.num\_entities(r)\}$  do
5:     get the entities expressed as a list of vertices  $(v_1, v_2, \dots) \leftarrow sort(c.get\_entity(r, o))$ 
6:     find entity in the map  $e \leftarrow VtE.find((v_1, v_2, \dots))$ 
7:     if the entity does not exist in the map already !j then
8:       store the entity in the map  $VtE.insert((v_1, v_2, \dots), ++i)$ 
9:       add information  $\mathcal{F}.EtC.append(\{c\}, \mathcal{F}.CtE[c].append(i))$ 
10:      else
11:       add information  $\mathcal{F}.EtC[e].append(c), \mathcal{F}.CtE[c].append(e)$ 
12:      end if
13:    end for
14:  end for
    
```




Fig. 39 **A** Vertex and edge ordinals (i.e., their ordering convention) for a triangular cell. **B** Vertex, edge, and cell indices for an example mesh consisting of three triangular elements

background edge, and the new vertex is a descendant of that edge.

4. The edge vertices descend from different background faces: the new vertex is inside the background element.
5. One edge vertex descends from a vertex and the other from a face or edge: the new vertex is a descendant of that face or edge.

B.2 Entity Connectivity

The algorithms presented require information about which lower-order entities, vertices, edges, and faces, are connected to cells and vice versa. This information is stored in the data structure shown in Fig. 10, and can be generated on any conformal mesh using Algorithm 12.

Example: edge-to-cell connectivity for a three-element mesh

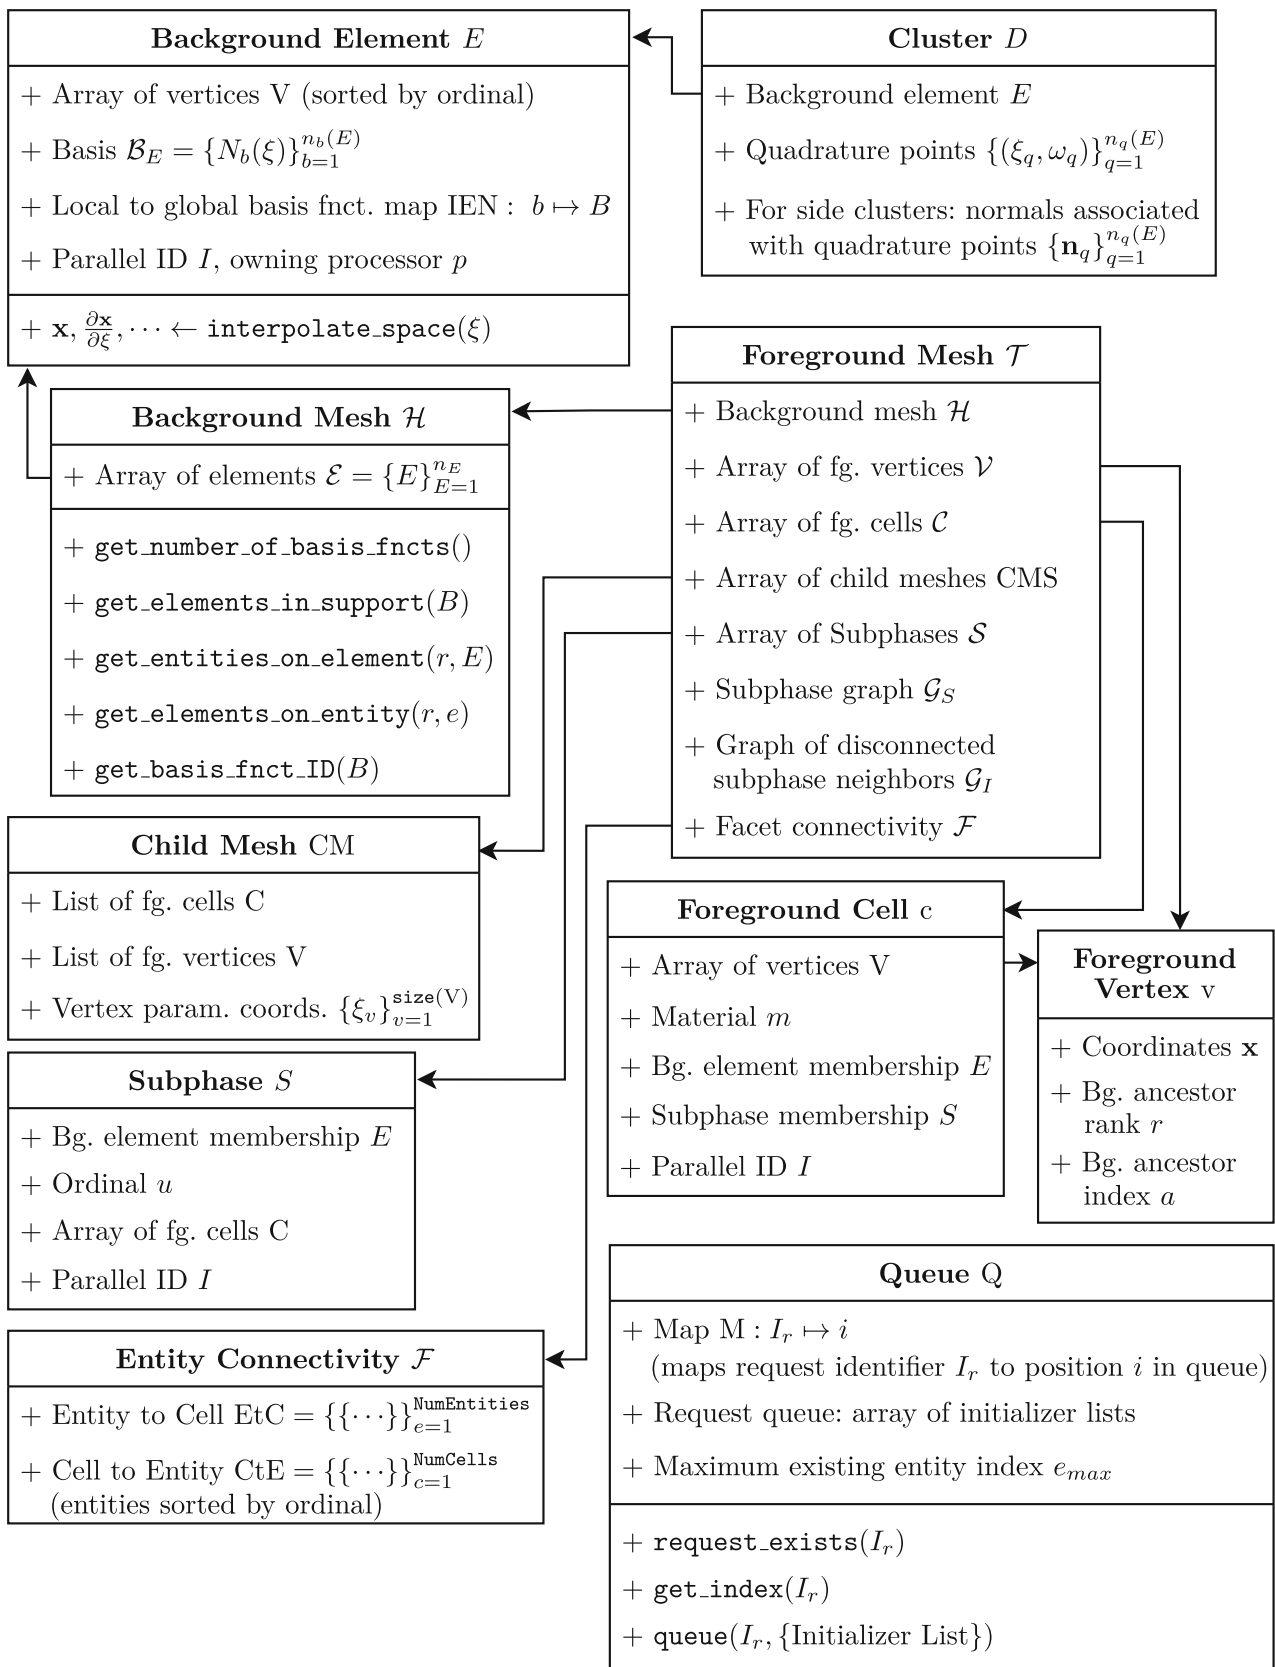


Fig. 40 UML chart of the data structures used in the presented preprocessor

Using the example previously shown in Fig. 10 in Section 3.1, we detail how to store and read/understand the entity-to-cell connectivity.

The cell connectivity \mathcal{C} declares the cells in the mesh by listing, in order, the vertices that define them. For the mesh shown in Fig. 39 (B) and using the ordering convention shown in Fig. 39 (A), the cell connectivity is:

$\mathcal{C} = \{$	cell connectivity
$\{1, 2, 3\},$	cell 1: vertices 1, 2, and 3
$\{4, 3, 2\},$	cell 2: vertices 4, 3, and 2
$\{3, 4, 1\}$	cell 3: vertices 3, 4, and 1

The entity-to-cell connectivity EtC for the edges, i.e., the edge-to-cell connectivity, lists the cells connected to each edge. For the example mesh in Fig. 39 (B), this is:

EtC = {	edge-to-cell connectivity
{1},	edge 1: connected to cell 1
{1, 2},	edge 2: connected to cells 1 and 2
{1, 3},	edge 3: connected to cells 1 and 3
{2, 3},	edge 4: connected to cells 2 and 3
{2},	edge 5: connected to cell 1
{3}}	edge 6: connected to cell 1

The cell-to-entity connectivity EtC for the edges, i.e., the cell-to-edge connectivity, lists the edges connected to each cell in the order of their ordinal. For the example mesh in Fig. 39 (B), this is:

CtE = {	cell-to-edge connectivity
$\{1, 2, 3\},$	cell 1: attached are edges 1, 2, and 3
$\{4, 2, 5\},$	cell 2: attached are edges 4, 2, and 5
$\{4, 6, 3\}$	cell 3: attached are edges 4, 6, and 3

B.3 Foreground Mesh – Complete Data Structure

The presented algorithms make use of data structures introduced in Figs. 8, 9, 13, and 17 throughout Section 3. The unified modeling language (UML) chart in Figure 40 illustrates their interaction with each other.

Acknowledgements N. Wunsch, J. A. Evans, and K. Maute were supported by the National Science Foundation (NSF), United States award OAC-2104106. N. Wunsch and K. Maute were additionally supported by the Air Force Office of Scientific Research (AFOSR), United States grant FA9550-20-1-0306. K. Doble was supported by Sandia National Laboratories through the contract PO2120843. M. R. Schmidt acknowledges the partial auspice of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344 (LLNL-JRNL-871539). L. Noël, M. R. Schmidt, J.A. Evans, and K.

Maute received support from the Defense Advanced Research Projects Agency (DARPA) under the TRADES program (agreement HR0011-17-2-0022). The geometry data for the example presented in Section 4.3 was provided by nTopology, Inc. free of charge. The opinions and conclusions presented in this paper are those of the authors and do not necessarily reflect the views of the sponsoring organizations.

Author Contributions The conceptual design and software development of the presented preprocessor were led by K. Doble and N. Wunsch. M. R. Schmidt and L. Noël have made significant contributions to both the design and development of the preprocessor. J. A. Evans and K. Maute supervised the presented work and contributed to the conceptual design. The numerical studies presented were performed by N. Wunsch. The first draft of the manuscript was written by N. Wunsch. All authors revised and commented on previous versions of the manuscript. All authors read and approved the final manuscript.

Data availability The numerical studies presented in this paper used the open-source software Moris, available at <https://github.com/kkmaute/moris>. The input decks for the presented example problems are archived in the same repository under /share/papers/XTK_paper. "LaTeX typeset, incl. hyperlinks, if this helps: " The numerical studies presented in this paper used the open-source software Moris, available at github.com/kkmaute/moris. The input decks for the presented example problems are archived in the same repository under /share/papers/XTK_paper."

Declarations

Conflict of interest J. A. Evans is an editor of Engineering with Computers. The other authors have no Conflict of interest to declare that are relevant to the content of this article.

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

- (2024) Moris – multi-physics optimization research and innovation system. <https://github.com/kkmaute/moris>, accessed: 2024-10-25
- Annarapu C, Hautefeuille M, Dolbow JE (2012) A robust Nitsche's formulation for interface problems. *Comput Methods Appl Mech Eng* 225:44–54. <https://doi.org/10.1016/j.cma.2012.03.008>
- Arndt D, Bangerth W, Bergbauer M et al (2023) The deal.II library, version 9.5. *J Numer Math* 31(3):231–246. <https://doi.org/10.1515/jnma-2023-0089>
- Arvo J, Kirk D (1989) A survey of ray tracing acceleration techniques. An Introduction to Ray Tracing, Academic Press, Ltd, pp 201–262
- Babuška I, Melenk JM (1997) The partition of unity method. *Int J Numer Methods Eng* 40(4):727–758

6. Badia S, Martín AF, Principe J (2018) Fempar: An object-oriented parallel finite element framework. *Arch Comput Methods Eng* 25(2):195–271
7. Badia S, Verdugo F, Martín AF (2018) The aggregated unfitted finite element method for elliptic problems. *Comput Methods Appl Mech Eng* 336:533–553. <https://doi.org/10.1016/j.cma.2018.03.022>
8. Badia S, Martorell PA, Verdugo F (2022) Geometrical discretisations for unfitted finite elements on explicit boundary representations. *J Comput Phys* 460:111162. <https://doi.org/10.1016/j.jcp.2022.111162>
9. Balay S, Gropp WD, McInnes LC, et al (1997) Efficient management of parallelism in object-oriented numerical software libraries. In: *Modern software tools for scientific computing*. Springer, p 163–202. https://doi.org/10.1007/978-1-4612-1986-6_8
10. Bazilevs Y, Calo VM, Cottrell JA et al (2010) Isogeometric analysis using T-splines. *Comput Methods Appl Mech Eng* 199(5–8):229–263. <https://doi.org/10.1016/j.cma.2009.02.036>
11. Belytschko T, Black T (1999) Elastic crack growth in finite elements with minimal remeshing. *Int J Numer Methods Eng* 45(5):601–620
12. Burman E (2010) Ghost penalty. *CR Math* 348(21–22):1217–1220. <https://doi.org/10.1016/j.crma.2010.10.006>
13. 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. <https://doi.org/10.1016/j.cma.2010.05.011>
14. 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. <https://doi.org/10.1016/j.apnum.2011.01.008>
15. Burman E, Hansbo P (2014) Fictitious domain methods using cut elements: III. A stabilized Nitsche method for Stokes' problem. *ESAIM Math Modelling Numer Anal-Modél Math Anal Numér* 48(3):859–874. <https://doi.org/10.1051/m2an/2013123>
16. Burman E, Claus S, Hansbo P et al (2015) Cutfem: discretizing geometry and partial differential equations. *Int J Numer Meth Eng* 104(7):472–501. <https://doi.org/10.1002/nme.4823>
17. Burman E, Hansbo P, Larson MG (2022) Cutfem based on extended finite element spaces. *Numer Math* 152(2):331–369. <https://doi.org/10.1007/s00211-022-01313-z>
18. Burman E, Hansbo P, Larson MG et al (2023) Extension operators for trimmed spline spaces. *Comput Methods Appl Mech Eng* 403:115707. <https://doi.org/10.1016/j.cma.2022.115707>
19. Claus S, Burman E, Massing A (2015) CutFEM: A stabilised Nitsche XFEM method for multi-physics problems. In: *23rd Conference on Computational Mechanics, ACME 2015*, 8–10 April 2015, Swansea, Wales, UK, Swansea University, pp 171–174
20. Coradello L, D'angella D, Carraturo M et al (2020) Hierarchically refined isogeometric analysis of trimmed shells. *Comput Mech* 66:431–447. <https://doi.org/10.1007/s00466-020-01858-6>
21. Cottrell JA, Hughes TJ, Bazilevs Y (2009) *Isogeometric analysis: toward integration of CAD and FEA*. John Wiley Sons. <https://doi.org/10.1002/9780470749081>
22. Duarte CA, Babuška I, Oden JT (2000) Generalized finite element methods for three-dimensional structural mechanics problems. *Comput Struct* 77(2):215–232. [https://doi.org/10.1016/S0045-7949\(99\)00211-4](https://doi.org/10.1016/S0045-7949(99)00211-4)
23. Duarte CA, Hamzeh O, Liszka T et al (2001) A generalized finite element method for the simulation of three-dimensional dynamic crack propagation. *Comput Methods Appl Mech Eng* 190(15–17):2227–2262. [https://doi.org/10.1016/S0045-7825\(00\)00233-4](https://doi.org/10.1016/S0045-7825(00)00233-4)
24. Düster A, Parvizian J, Yang Z et al (2008) The finite cell method for three-dimensional problems of solid mechanics. *Comput Methods Appl Mech Eng* 197(45–48):3768–3782. <https://doi.org/10.1016/j.cma.2008.02.036>
25. Elfverson D, Larson MG, Larsson K (2018) CutIGA with basis function removal. *Adv Model Simul Eng Sci* 5(1):1–19. <https://doi.org/10.1186/s40323-018-0099-2>
26. Elhaddad M, Zander N, Bog T et al (2018) Multi-level hp-finite cell method for embedded interface problems with application in biomechanics. *Int J Numer Methods Biomed Eng* 34(4):e2951. <https://doi.org/10.1002/cnm.2951>
27. Fernández-Méndez S, Huerta A (2004) Imposing essential boundary conditions in mesh-free methods. *Comput Methods Appl Mech Eng* 193(12–14):1257–1275. <https://doi.org/10.1016/j.cma.2003.12.019>
28. Fries TP, 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. <https://doi.org/10.1002/nme.2914>
29. Fries TP, Omerović S (2016) Higher-order accurate integration of implicit geometries. *Int J Numer Meth Eng* 106(5):323–371. <https://doi.org/10.1002/nme.5121>
30. Fromm JE, Wunsch N, Xiang R et al (2023) Interpolation-based immersed finite element and isogeometric analysis. *Comput Methods Appl Mech Eng* 405:115890. <https://doi.org/10.1016/j.cma.2023.115890>
31. Fromm JE, Wunsch N, Maute K et al (2024) Interpolation-based immersogeometric analysis methods for multi-material and multi-physics problems. *Comput Mech* 75:1–25. <https://doi.org/10.1007/s00466-024-02506-z>
32. Gabriel E, Fagg GE, Bosilca G, et al (2004) Open mpi: Goals, concept, and design of a next generation mpi implementation. In: *Recent Advances in Parallel Virtual Machine and Message Passing Interface: 11th European PVM/MPI Users' Group Meeting Budapest, Hungary, September 19–22, 2004*. Proceedings 11, Springer, pp 97–104. https://doi.org/10.1007/978-3-540-30218-6_19
33. Groen JP, Langelaar M, Sigmund O et al (2017) Higher-order multi-resolution topology optimization using the finite cell method. *Int J Numer Meth Eng* 110(10):903–920. <https://doi.org/10.1002/nme.5432>
34. Gunderman D, Weiss K, Evans JA (2021) High-accuracy mesh-free quadrature for trimmed parametric surfaces and volumes. *Comput Aided Des* 141:103093. <https://doi.org/10.1016/j.cad.2021.103093>
35. Gunderman D, Weiss K, Evans JA (2021) Spectral mesh-free quadrature for planar regions bounded by rational parametric curves. *Comput Aided Des* 130:102944. <https://doi.org/10.1016/j.cad.2020.102944>
36. 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. [https://doi.org/10.1016/S0045-7825\(02\)00524-8](https://doi.org/10.1016/S0045-7825(02)00524-8)
37. Hansbo A, Hansbo P (2004) A finite element method for the simulation of strong and weak discontinuities in solid mechanics. *Comput Methods Appl Mech Eng* 193(33–35):3523–3540. <https://doi.org/10.1016/j.cma.2003.12.041>
38. Heimann F, Lehrenfeld C, Preuß J (2023) Geometrically higher order unfitted space-time methods for pdes on moving domains. *SIAM J Sci Comput* 45(2):B139–B165. <https://doi.org/10.1137/22M1476034>
39. Hernández V, Román JE, Tomás A, et al (2007) Krylov-Schur methods in SLEPc. *Universitat Politècnica de Valencia, Tech Rep STR-7*
40. Hsu MC, Kamensky D, Bazilevs Y et al (2014) Fluid-structure interaction analysis of bioprosthetic heart valves: significance of arterial wall deformation. *Comput Mech* 54:1055–1071. <https://doi.org/10.1007/s00466-014-1059-4>
41. Hsu MC, Kamensky D, Xu F et al (2015) Dynamic and fluid-structure interaction simulations of bioprosthetic heart valves using parametric design with T-splines and fung-type material mod-

- els. *Comput Mech* 55:1211–1225. <https://doi.org/10.1007/s00466-015-1166-x>
42. Hughes TJ (2012) The finite element method: linear static and dynamic finite element analysis. Courier Corporation
 43. Jenkins N, Maute K (2016) An immersed boundary approach for shape and topology optimization of stationary fluid-structure interaction problems. *Struct Multidiscip Optim* 54:1191–1208. <https://doi.org/10.1007/s00158-015-1229-9>
 44. Jiang W, Spencer BW, Dolbow JE (2020) Ceramic nuclear fuel fracture modeling with the extended finite element method. *Eng Fract Mech* 223:106713. <https://doi.org/10.1016/j.engfractmech.2019.106713>
 45. Johansson A, Kehlet B, Larson MG et al (2019) Multimesh finite element methods: Solving PDEs on multiple intersecting meshes. *Comput Methods Appl Mech Eng* 343:672–689. <https://doi.org/10.1016/j.cma.2018.09.009>
 46. Kamensky D, Hsu MC, Schillinger D et al (2015) An immersed-geometric variational framework for fluid-structure interaction: Application to bioprosthetic heart valves. *Comput Methods Appl Mech Eng* 284:1005–1053. <https://doi.org/10.1016/j.cma.2014.10.040>
 47. Lehrenfeld C (2016) High order unfitted finite element methods on level set domains using isoparametric mappings. *Comput Methods Appl Mech Eng* 300:716–733. <https://doi.org/10.1016/j.cma.2015.12.005>
 48. Lehrenfeld C, Heimann F, Preuß J et al (2021) ngsxfem: Add-on to ngsolve for geometrically unfitted finite element discretizations. *J Open Source Softw* 6(64):3237. <https://doi.org/10.21105/joss.03237>
 49. Lian Z, Leidinger L, Hartmann S et al (2025) Cross-talk effects in trimmed isogeometric shells and the control point duplication approach. *Comput Methods Appl Mech Eng* 438:117849. <https://doi.org/10.1016/j.cma.2025.117849>
 50. 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. <https://doi.org/10.1016/j.jcp.2017.10.026>
 51. 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. <https://doi.org/10.1016/j.jcp.2018.01.023>
 52. Martorell PA, Badia S (2024) High order unfitted finite element discretizations for explicit boundary representations. *J Comput Phys* 511:113127. <https://doi.org/10.1016/j.jcp.2024.113127>
 53. McCalpin JD et al (1995) Memory bandwidth and machine balance in current high performance computers. *IEEE Comput Soc Tech Comm Comput Archit (TCCA) News* 2:19–25
 54. 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. [https://doi.org/10.1016/S0045-7825\(96\)01087-0](https://doi.org/10.1016/S0045-7825(96)01087-0)
 55. Moës N, Dolbow J, Belytschko T (1999) A finite element method for crack growth without remeshing. *Int J Numer Methods Eng* 46(1):131–150
 56. Mourrain B, Rouillier F, Roy MF (2004) Bernstein’s basis and real root isolation. PhD thesis, INRIA, France
 57. 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. <https://doi.org/10.1002/nme.4569>
 58. Nagai T (2018) Space-time extended finite element method with applications to fluid-structure interaction problems. PhD thesis, University of Colorado at Boulder, US
 59. Nitsche J (1971) Über ein variationsprinzip zur lösung von dirichlet-problemen bei verwendung von teilträumen, die keinen randbedingungen unterworfen sind. In: *Abhandlungen aus dem mathematischen Seminar der Universität Hamburg*, Springer, pp 9–15. <https://doi.org/10.1007/BF02995904>
 60. Noël L, Duysinx P (2017) Shape optimization of microstructural designs subject to local stress constraints within an XFEM-level set framework. *Struct Multidiscip Optim* 55:2323–2338. <https://doi.org/10.1007/s00158-016-1642-8>
 61. Noël L, Maute K (2023) XFEM level set-based topology optimization for turbulent conjugate heat transfer problems. *Struct Multidiscip Optim* 66(1):2. <https://doi.org/10.1007/s00158-022-03353-3>
 62. Noël L, Schmidt M, Doble K et al (2022) XIGA: An extended isogeometric analysis approach for multi-material problems. *Comput Mech* 70(6):1281–1308. <https://doi.org/10.1007/s00466-022-02200-y>
 63. Parvzian J, Düster A, Rank E (2007) Finite cell method: h-and p-extension for embedded domain problems in solid mechanics. *Comput Mech* 41(1):121–133. <https://doi.org/10.1007/s00466-007-0173-y>
 64. 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. <https://doi.org/10.1016/j.camwa.2018.03.032>
 65. de Prenter F, Verhoosel CV, van Brummelen EH et al (2023) Stability and conditioning of immersed finite element methods: analysis and remedies. *Archiv Comput Methods Eng*. <https://doi.org/10.1007/s11831-023-09913-0>
 66. Reuter M, Mikkelsen TS, Sherbrooke EC et al (2008) Solving non-linear polynomial systems in the barycentric Bernstein basis. *Vis Comput* 24:187–200. <https://doi.org/10.1007/s00371-007-0184-x>
 67. Ruess M, Tal D, Trabelsi N et al (2012) The finite cell method for bone simulations: verification and validation. *Biomech Model Mechanobiol* 11:425–437. <https://doi.org/10.1007/s10237-011-0322-2>
 68. Saad Y (1993) A flexible inner-outer preconditioned gmres algorithm. *SIAM J Sci Comput* 14(2):461–469. <https://doi.org/10.1137/0914028>
 69. Sauerland H, Fries TP (2013) The stable XFEM for two-phase flows. *Comput Fluids* 87:41–49. <https://doi.org/10.1016/j.compfluid.2012.10.017>
 70. Saye RI (2022) High-order quadrature on multi-component domains implicitly defined by multivariate polynomials. *J Comput Phys* 448:110720. <https://doi.org/10.1016/j.jcp.2021.110720>
 71. Schenk O, Gärtner K (2004) Solving unsymmetric sparse systems of linear equations with pardiso. *Futur Gener Comput Syst* 20(3):475–487. <https://doi.org/10.1016/j.future.2003.07.011>
 72. Schenk O, Gärtner K (2006) On fast factorization pivoting methods for sparse symmetric indefinite systems. *Electron Trans Numer Anal* 23(1):158–179
 73. 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. <https://doi.org/10.1007/s11831-014-9115-y>
 74. Schillinger D, Harari I, Hsu MC et al (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. <https://doi.org/10.1016/j.cma.2016.06.026>
 75. Schmidt M, Noël L, Doble K et al (2023) Extended isogeometric analysis of multi-material and multi-physics problems using hierarchical B-splines. *Comput Mech* 71(6):1179–1203. <https://doi.org/10.1007/s00466-023-02306-x>
 76. Scholz F, Jüttler B (2019) Numerical integration on trimmed three-dimensional domains with implicitly defined trimming surfaces. *Comput Methods Appl Mech Eng* 357:112577. <https://doi.org/10.1016/j.cma.2019.112577>

77. Scholz F, Jüttler B (2020) High-order quadrature on planar domains based on transport theorems for implicitly defined moving curves. *Tech. rep.*, Tech. Rep. 89, Nationales Forschungsnetzwerk, Geometry+ Simulation (May 2020)
78. Schott B, Wall W (2014) A new face-oriented stabilized XFEM approach for 2d and 3d incompressible Navier-Stokes equations. *Comput Methods Appl Mech Eng* 276:233–265. <https://doi.org/10.1016/j.cma.2014.02.014>
79. Schweitzer MA (2012) Generalizations of the finite element method. *Central Eur J Math* 10:3–24. <https://doi.org/10.2478/s11533-011-0112-1>
80. Shalf J (2020) The future of computing beyond Moore's Law. *Philos Trans R Soc A Math Phys Eng Sci* 378(2166):20190061. <https://doi.org/10.1098/rsta.2019.0061>
81. Soghrati S (2014) Hierarchical interface-enriched finite element method: an automated technique for mesh-independent simulations. *J Comput Phys* 275:41–52. <https://doi.org/10.1016/j.jcp.2014.06.016>
82. Soghrati S, Geubelle PH (2012) A 3D interface-enriched generalized finite element method for weakly discontinuous problems with complex internal geometries. *Comput Methods Appl Mech Eng* 217:46–57. <https://doi.org/10.1016/j.cma.2011.12.010>
83. Soghrati S, Aragón AM, Armando Duarte C et al (2012) An interface-enriched generalized fem for problems with discontinuous gradient fields. *Int J Numer Meth Eng* 89(8):991–1008. <https://doi.org/10.1002/nme.3273>
84. Stanford JW, Fries TP (2019) A higher-order conformal decomposition finite element method for plane B-rep geometries. *Comput Struct* 214:15–27. <https://doi.org/10.1016/j.compstruc.2018.12.006>
85. Stewart GW (2002) A Krylov-Schur algorithm for large eigenproblems. *SIAM J Matrix Anal Appl* 23(3):601–614. <https://doi.org/10.1137/S0895479800371529>
86. Sudhakar Y, Wall WA (2013) Quadrature schemes for arbitrary convex/concave volumes and integration of weak form in enriched partition of unity methods. *Comput Methods Appl Mech Eng* 258:39–54. <https://doi.org/10.1016/j.cma.2013.01.007>
87. Terada K, Asai M, Yamagishi M (2003) Finite cover method for linear and non-linear analyses of heterogeneous solids. *Int J Numer Meth Eng* 58(9):1321–1346. <https://doi.org/10.1002/nme.820>
88. Van Dijk NP, Maute K, Langelaar M et al (2013) Level-set methods for structural topology optimization: a review. *Struct Multidiscip Optim* 48:437–472. <https://doi.org/10.1007/s00158-013-0912-y>
89. Verhoosel CV, Van Zwieten G, Van Rietbergen B et al (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. <https://doi.org/10.1016/j.cma.2014.07.009>
90. Villanueva CH, Maute K (2017) CutFEM topology optimization of 3d laminar incompressible flow problems. *Comput Methods Appl Mech Eng* 320:444–473. <https://doi.org/10.1016/j.cma.2017.03.007>
91. Vlahinos M, O'Hara R (2020) Unlocking advanced heat exchanger design and simulation with nTop platform and ANSYS CFX. *NTopology Inc* 8
92. Wade C, Williams G, Connelly S et al (2023) OpenVCAD: An open source volumetric multi-material geometry compiler. *Addit Manuf* 79(103912):103912. <https://doi.org/10.1016/j.addma.2023.103912>
93. Zhang J, Zhebel E, van den Boom SJ et al (2022) An object-oriented geometric engine design for discontinuities in unfitted/immersed/enriched finite element methods. *Int J Numer Meth Eng* 123(21):5126–5154. <https://doi.org/10.1002/nme.7049>
94. Zhang L, Gerstenberger A, Wang X et al (2004) Immersed finite element method. *Comput Methods Appl Mech Eng* 193(21–22):2051–2067. <https://doi.org/10.1137/1.9780898717464.ch8>
95. Zhang Z, Jiang W, Dolbow JE et al (2018) A modified moment-fitted integration scheme for X-FEM applications with history-dependent material data. *Comput Mech* 62:233–252. <https://doi.org/10.1007/s00466-018-1544-2>
96. Zhu T, Atluri S (1998) A modified collocation method and a penalty formulation for enforcing the essential boundary conditions in the element free galerkin method. *Comput Mech* 21(3):211–222. <https://doi.org/10.1007/s004660050296>

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