



LAWRENCE
LIVERMORE
NATIONAL
LABORATORY

LLNL-JRNL-842737

Extended isogeometric analysis of multi-material and multi-physics problems using hierarchical B-splines

M. R. Schmidt, L. Noel, K. Doble, J. A. Evans, K. Maute

November 29, 2022

Computational Mechanics

Disclaimer

This document was prepared as an account of work sponsored by an agency of the United States government. Neither the United States government nor Lawrence Livermore National Security, LLC, nor any of their employees makes any warranty, expressed or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States government or Lawrence Livermore National Security, LLC. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States government or Lawrence Livermore National Security, LLC, and shall not be used for advertising or product endorsement purposes.

Extended isogeometric analysis of multi-material and multi-physics problems using hierarchical B-splines

Mathias Schmidt¹, Lise Noël², Keenan Doble³, John A. Evans³, Kurt Maute³

Received: date / Accepted: date

Abstract This paper presents an immersed, isogeometric finite element framework to predict the response of multi-material, multi-physics problems with complex geometries using locally refined discretizations. To circumvent the need to generate conformal meshes, this work uses an eXtended Finite Element Method (XFEM) to discretize the governing equations on non-conforming, embedding meshes. A flexible approach to create truncated hierarchical B-splines discretizations is presented. This approach enables the refinement of each state variable field individually to meet field-specific accuracy requirements. To obtain an immersed geometry representation that is consistent across all hierarchically refined B-spline discretizations, the geometry is immersed into a single mesh, the XFEM background mesh, which is constructed from the union of all hierarchical B-spline meshes. An extraction operator is introduced to represent the truncated hierarchical B-spline bases in terms of Lagrange shape functions on the XFEM background mesh without loss of accuracy. The truncated hierarchical B-spline bases are enriched using a generalized Heaviside enrichment strategy to accommodate small geometric features and multi-material problems. The governing equations are augmented by a formulation of the face-oriented ghost stabilization enhanced for locally refined B-spline bases. We present examples for

two- and three-dimensional linear elastic and thermo-elastic problems. The numerical results validate the accuracy of our framework. The results also demonstrate the applicability of the proposed framework to large, geometrically complex problems.

Keywords Immersed Finite Element Method, eXtended Isogeometric Analysis, Multi-material Problems, Multi-physics Problems, Truncated Hierarchical B-splines, Lagrange Extraction

1 Introduction

Finite element analysis is frequently used to predict the response of systems described by partial differential equations defined over a spatial domain. In classical finite element methods, the domain is discretized using a single mesh that conforms to the external boundaries and internal material interfaces. For problems with complex shapes and multiple material phases, the construction of this conformal mesh is often a major bottleneck in the analysis process, see Bazilevs et al. (2010). Conformal mesh generation may also hamper the automation of the finite element analysis for problems with changing geometry as encountered in, for example, phase-change problems with dynamically evolving interfaces or shape and topology optimization.

The state variable fields may exhibit large spatial gradients at boundaries and material interfaces, as well as in the vicinity of small geometric features. To resolve these spatial gradients, a sufficiently fine discretization is needed. While uniformly fine discretization may yield an accurate approximation, the associated computational cost may exceed practical limits. A locally refined mesh balances discretization needs and computational cost. However, generating locally refined meshes

¹Computational Engineering Division, Lawrence Livermore National Laboratory, 7000 East Ave, Livermore, 94550, California, USA

²Department of Precision and Microsystems Engineering, Faculty of Mechanical, Maritime and Materials Engineering, Delft University of Technology, Mekelweg 2, 2628 CD Delft, The Netherlands

³Aerospace Mechanics Research Center, Department of Aerospace Engineering Sciences, University of Colorado Boulder, 3775 Discovery Dr, Boulder, 80309-0429, Colorado, USA
E-mail: schmidt43@llnl.gov

for standard finite element methods further increases the complexity of mesh generation.

Many engineering and science applications involve multiple, often coupled state variable fields. For such problems, each field may require differently refined discretizations. For example, consider the two-material, thermo-elastic problem of a circle embedded in a rectangular plate shown in Fig. 1a. We assume that the circle is made of a material with a finite coefficient of thermal expansion (CTE), while the plate's CTE is zero. The plate is subject to a spatially varying heat flux along its right edge and clamped at its left edge. The heat flux increases the temperature in the system, causing the circle to expand which generates stresses in both the plate and circle due to the CTE mismatch. To predict the stress field at the material interface with high accuracy, accurate approximations of both the temperature and displacement fields are required. The contour of the norm of the diffusive flux is shown in Fig. 1b and the contour of the Von Mises stress field in Fig. 1c. The regions with large spatial gradients differ between the displacement and the temperature field. Separately adapting the discretizations associated with each field enables a sufficient resolution of each field while minimizing the overall computational cost. An example of such discretizations for the temperature and displacement fields is depicted in Fig. 2a and 2b, respectively.

Lagrange polynomials are the standard choice for basis functions in finite elements. The element-local nature of this class of basis functions and their interpolatory property facilitate a convenient and efficient implementation and enforcement of Dirichlet boundary and interface conditions. However, Lagrange basis functions limit the inter-element continuity to C^0 , irrespective of the polynomial order of the approximation within the element. The lack of higher-order inter-element continuity affects the efficiency of Lagrange bases, measured by the number of degrees of freedom needed to achieve a desired accuracy; see for example Evans et al. (2009).

In this work, we propose a computational framework that addresses the issues encountered with standard finite element approaches as outlined above. This framework synthesizes the following concepts which will be further discussed in detail below. An immersed finite element approach is adopted to eliminate the need for conformal mesh generation. Higher order B-spline basis functions are used to increase the discretization efficiency over Lagrange bases. For each state variable, a truncated hierarchically refined B-spline basis is generated to meet field-specific accuracy requirements.

Immersed Boundary Methods (IBMs) have gained in popularity in recent years, see Babuška and Melenk (1997); Peskin (2002); Mittal and Iaccarino (2005). The

general idea of these methods is to immerse the geometry of the physical domain into a computational domain with a much simpler geometry. Thus, the geometries of the physical and computational domains are decoupled. The geometric simplicity of the computational domain allows for the convenient generation of IBM background meshes, such as tensor grids. The geometry is immersed into the IBM background mesh, which simplifies the discretization of the state variable fields.

In this work, we focus on finite element formulations of IBMs. Immersed Finite Element Methods (IFEMs), also referred to as geometrically unfitted or embedded domain finite element methods, include the Finite Cell Method , see Parvizian et al. (2007); Düster et al. (2008); Schillinger and Ruess (2015), the CutFEM (Remacle et al. (2012)), the Generalized Finite Element Method (GFEM) (Duarte et al. (2000); Strouboulis et al. (2000)), and the eXtended Finite Element Method (XFEM), see Belytschko and Black (1999); Belytschko et al. (2009). In this paper, we consider specifically the latter approach. The XFEM augments the standard finite element basis with additional basis functions to represent discontinuities of the state variable field within an XFEM background element that is intersected by a boundary or an interface. The augmented finite element basis satisfies the partition of unity (PU) property. In this work, we adopt a Heaviside enrichment strategy for its flexibility in modeling interface and boundary conditions of multi-material problems with complex geometries, see Noël et al. (2022).

Traditionally, Heaviside enriched XFEM approaches discretize state variable fields by Lagrange basis functions which are defined on the XFEM background mesh. In this paper, we adopt higher order B-splines for discretizing state variable fields. The advantages of B-splines basis functions for finite element methods have been demonstrated in the context of Isogeometric Analysis (IGA), see Hughes et al. (2005). While IGA was originally developed to eliminate the discrepancy between CAD geometry representation and finite element analysis, Evans et al. (2009) showed that B-spline basis functions in general improve accuracy, robustness, and computational efficiency. In this work, we further utilize the refineability property of B-splines, see Garau and Vázquez (2018).

B-spline basis functions were studied with IFEMs by Schillinger et al. (2012); Schillinger and Rank (2011); Schillinger and Ruess (2015); Verhoosel et al. (2015); Elfverson et al. (2018); Divi et al. (2020); Noël et al. (2020). Nguyen (2012) introduced B-spline approximations to the XFEM and coined the term *X-IGA*. These works demonstrated that integrating higher order B-splines discretizations into IFEM approaches yields im-

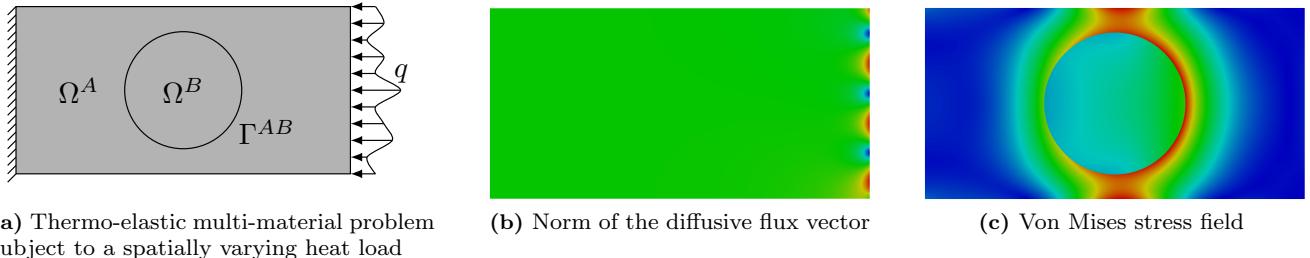


Fig. 1: Illustration of an example thermo-elastic problem

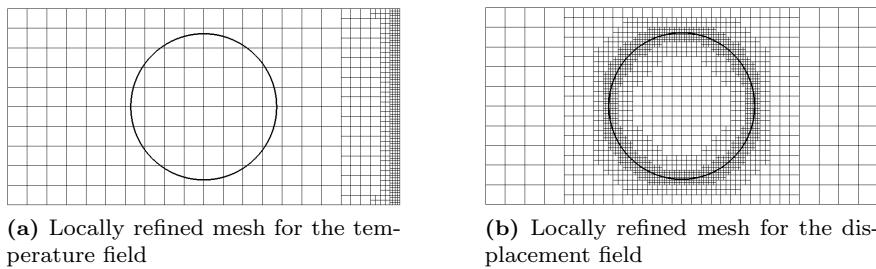


Fig. 2: Example of locally refined discretizations for the thermo-elastic problem depicted in Fig. 1a

proved accuracy and computational efficiency compared to the use of Lagrange basis functions.

The ability to locally refine B-splines enables the convenient construction of locally refined approximation spaces, see Giannelli et al. (2012); Buffa and Giannelli (2017); Bracco et al. (2019). While standard hierarchical B-spline bases do not fulfill the PU property, truncating the bases restores the PU property Giannelli et al. (2012). Truncated Hierarchical B-spline (THB) basis functions form a sparse, strongly stable basis Giannelli et al. (2012) and are employed in this work.

This paper contributes to the work on IFEMs as follows. We introduce a versatile discretization approach based on THBs defined on locally refined tensor meshes. Each state variable can be discretized individually by choosing the polynomial order and the local refinement independently from the discretization of other fields. This approach leads to an explicit control over the accuracy of the field approximations and the computational cost of the analysis. Computationally efficient algorithms and data structures are introduced to enable the application of this discretization approach to large-scale problems in two and three dimensions using parallel computing. To integrate the locally refined B-spline discretization approach into the XFEM, a single union mesh is generated from the individual B-spline meshes. Fig. 3 shows the union mesh generated for the two locally refined meshes presented in Fig. 2. The union mesh serves as the XFEM background mesh in which the geometry is immersed. To increase the geometry

resolution, the XFEM background mesh can be additionally refined. We introduce an exact extraction operator that expresses THBs in terms of Lagrange basis functions over each element of the XFEM background mesh.

The enrichment strategy of Makhija and Maute (2014); Villanueva and Maute (2014); Noël et al. (2022) is generalized to enrich the locally refined THB basis functions considering their support for a given intersection geometry. The XFEM problem is augmented by an X-IGA formulation of the face-oriented ghost stabilization from Noël et al. (2022). The governing equations are integrated by standard quadrature rules on a geometry-conforming integration mesh that is constructed by cutting XFEM background elements that are intersected by the external boundaries or internal material interfaces. Boundary and interface conditions are enforced weakly by Nitsche's method. We illustrate the main characteristics and features of the proposed XFEM framework by numerical examples considering linear elastic and thermo-elastic problems. We study the convergence of geometric and state variable discretization errors with uniform and local mesh refinement for different B-spline orders. To demonstrate the applicability of our XFEM framework to complex multi-material, multi-physics problems, we apply our XFEM framework to the thermo-elastic analysis of a 3D polycrystalline micro-structure.

This paper is organized as follows: Section 2 outlines the Heaviside-enriched XFEM framework. Section

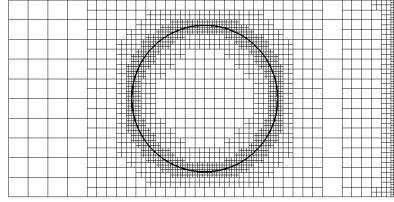


Fig. 3: Union of the locally refined meshes depicted in Figures 2a and 2b

3 recalls the fundamentals of hierarchical B-splines as well as their truncation to restore the PU property. Section 4 details the meshing algorithms and data structures. Sections 5 summarizes the governing equations, Nitsche's formulation, and face-oriented ghost stabilization for linear thermo-elasticity. Numerical two and three dimensional examples are presented in Section 6. The main findings are summarized in Section 7, together with recommendations for future work.

2 The eXtended Finite Element Method

In this work, we adopt the XFEM to perform analysis on non-conforming background meshes. We follow the basic concepts of a generalized Heaviside enrichment strategy, introduced by Terada et al. (2003). In this section, the basic concepts of the XFEM as relevant for this work are briefly described.

We start from a non-conforming finite element approximation space with THB basis functions $\{B_j\}_{j=1}^{\tilde{n}}$ that are defined on a locally refined discretization. To approximate the state variable fields in different material phases, each basis function is enriched. The enrichment of a particular basis function depends on the number of topologically disconnected regions of all material phases in the support of this basis function. The approach is illustrated in Fig. 4 for a configuration with two material phases, Ω^1 and Ω^2 . The support of the basis function B is depicted by the dashed red line. It spans three topologically disconnected regions, each occupied by one of the two material phases. Therefore, the basis function is enriched three times.

This approach is applicable to any number of material phases, intersection configurations, and basis function support. To define the enriched basis, we introduce the indicator function ψ_j^ℓ . This function is equal to 1 for points located within the portion of the support of the j^{th} basis function corresponding to the enrichment level ℓ and equal to zero elsewhere. The enriched basis then is simply $\{B_j \psi_j^\ell : j \in \{1, \dots, \tilde{n}\} \text{ and } \ell \in \{1, \dots, L_j\}\}$ where L_j is the maximum number of enrichment levels for basis function B_j . Since the original THB basis satisfies the PU principle, so does the enriched THB basis.

The enriched finite element approximation of a vector-valued state variable $\mathbf{u}^h(\mathbf{x})$ can be written in terms of the enriched THB basis as follows:

$$\mathbf{u}^h(\mathbf{x}) = \sum_{j=1}^{\tilde{n}} \left(\sum_{l=1}^{L_j} \psi_j^\ell(\mathbf{x}) B_j(\mathbf{x}) \mathbf{c}_j^\ell \right), \quad (1)$$

where \mathbf{c}_j^ℓ is the coefficient associated to the j^{th} original THB basis function and the j^{th} enrichment level. The indicator function ψ_j^ℓ enforces that only one set of enriched basis functions is used to describe the state variable at each point in the domain. A more detailed description of the enrichment strategy can be found in Noël et al. (2022).

The Heaviside enriched XFEM formulation outlined above enables the modeling of C^{-1} intra-element discontinuities of state variables within a non-conforming background element. Essential boundary conditions can be enforced weakly by, for example, Nitsche's method (Nitsche (1971); Burman (2012)) or the stabilized Lagrange multiplier method (Gerstenberger and Wall (2008)).

Immersing geometry into the XFEM background mesh can result in basis functions with small support within the geometric domain, leading to poorly conditioned systems of discretized governing equations. Various strategies to mitigate this issue have been studied in the literature, such as the face-oriented ghost stabilization (Burman (2010); Burman and Hansbo (2014); Noël et al. (2022)), basis function removal (Embar et al. (2010); Elfverson et al. (2018)), and pre-conditioning (Lang et al. (2014); de Prenter et al. (2017)). In this work, we extend the face-oriented ghost stabilization to hierarchically refined B-spline discretizations; as further described in Subsection 5.2.

3 Hierarchical B-splines

This section focuses on hierarchical B-splines for a locally refined discretization of state variable fields. First, the basic concepts of B-splines in one and multiple dimensions are recalled. Then the B-spline refinement and the construction of non-truncated (HB) and truncated hierarchical B-spline (THB) bases are described.

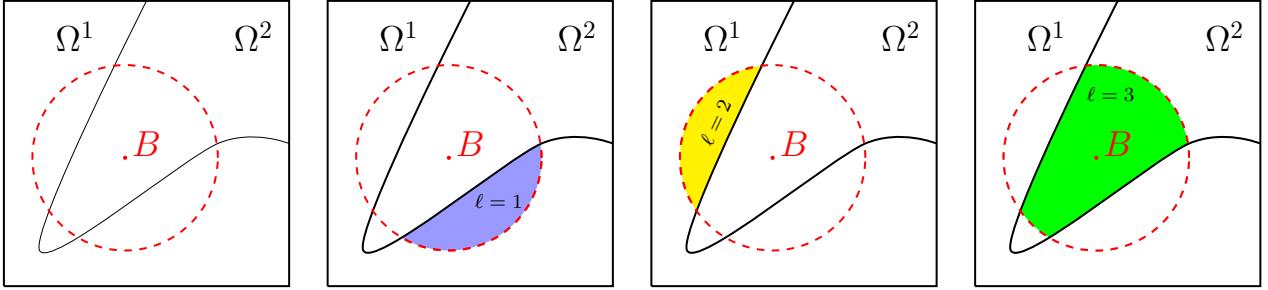


Fig. 4: Basis function enrichment for a two-material problem

3.1 B-spline Basis Functions

In 1D, we define a knot vector $\Xi = \{\xi_1, \xi_2, \dots, \xi_{n+p+1}\}$, for which $\xi \in \mathbb{R}$ and $\xi_1 \leq \xi_2 \leq \dots \leq \xi_{n+p+1}$. A univariate B-spline basis function $N_{i,p}(\xi)$ of degree p is constructed recursively starting from the piecewise constant basis function:

$$N_{i,0}(\xi) = \begin{cases} 1, & \text{if } \xi_i \leq \xi \leq \xi_{i+1}, \\ 0, & \text{otherwise,} \end{cases} \quad (2)$$

and using the Cox de Boor recursion formula (de Boor (1972)) for higher degrees, $p > 0$:

$$N_{i,p}(\xi) = \frac{\xi - \xi_i}{\xi_{i+p} - \xi_i} N_{i,p-1}(\xi) + \frac{\xi_{i+p+1} - \xi}{\xi_{i+p+1} - \xi_{i+1}} N_{i+1,p-1}(\xi). \quad (3)$$

To guarantee a C^{p-1} continuity over the entire computational domain, none of the interior knots should be repeated. The corresponding B-spline basis exhibits a C^{p-1} continuity at every knot in the interior of the domain, while it is C^∞ continuous in between the knots. A knot span is defined as the half open interval $[\xi_i, \xi_{i+1})$ and a B-spline element is defined as a non-empty knot span.

In 2D and 3D, tensor-product B-spline basis functions $B_{i,p}(\xi)$ are constructed by applying the tensor-product operation to univariate B-spline basis functions in each parametric direction. Denoting the parametric space dimension as d_p , a tensor-product B-spline basis is constructed starting from d_p knot vectors $\Xi^m = \{\xi_1^m, \xi_2^m, \dots, \xi_{n_m+p_m+1}^m\}$ with p_m being the polynomial degree and n_m the number of basis functions in the parametric direction $m = 1, \dots, d_p$. A tensor-product B-spline basis function is generated from d_p univariate B-spline basis functions $N_{i_m,p_m}^m(\xi^m)$ in each parametric direction m as follows:

$$B_{\mathbf{i},\mathbf{p}}(\xi) = \prod_{m=1}^{d_p} N_{i_m,p_m}^m(\xi^m), \quad (4)$$

where the position in the tensor-product structure is given by the index vector $\mathbf{i} = \{i_1, \dots, i_{d_p}\}$, and the vector $\mathbf{p} = \{p_1, \dots, p_{d_p}\}$ defines the polynomial degree in each direction. Similarly to the univariate case, an element is defined as the tensor-product of d_p non-empty knot spans. Additionally, a B-spline space \mathcal{V} is defined as the span of B-spline basis functions.

3.2 B-spline Refinement

Hierarchical refinement of uniform B-splines is achieved by subdivision. A univariate B-spline basis function is expressed as a linear combination of $p+2$ contracted, translated, and scaled copies of itself:

$$N_p(\xi) = 2^{-p} \sum_{j=0}^{p+1} \binom{p+1}{j} N_p(2\xi - j), \quad (5)$$

where the binomial coefficient is defined as:

$$\binom{p+1}{j} = \frac{(p+1)!}{j!(p+1-j)!}. \quad (6)$$

Fig. 5 shows the refinement of a quadratic univariate B-spline basis function obtained by subdivision.

The subdivision in Eq. (4) for a univariate B-spline basis can be extended to tensor-product B-spline basis functions B_p as follows, see Schillinger et al. (2012):

$$B_{\mathbf{p}}(\xi) = \sum_{\mathbf{j}} \left(\prod_{m=1}^d 2^{-p_m} \binom{p_m+1}{j_m} N_{p_m}(2\xi^m - j_m) \right), \quad (7)$$

where the index vector $\mathbf{j} = \{i_1, \dots, i_{d_p}\}$ collects the positions in the tensor-product structure.

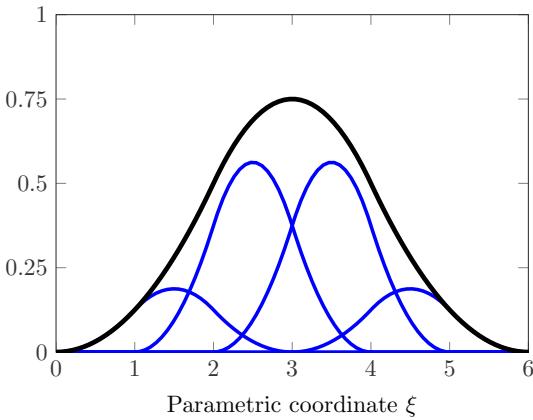


Fig. 5: Subdivision of a quadratic B-spline basis function (black) into $p + 2$ contracted B-spline basis functions (blue) of half the knot span width

3.3 Hierarchical B-splines

To define a hierarchical mesh of depth n , a sequence of subdomains Ω^l is introduced:

$$\Omega^{n-1} \subseteq \Omega^{n-2} \subseteq \dots \subseteq \Omega^0 = \Omega, \quad (8)$$

where each subdomain Ω^l is a refined sub-region of Ω^{l-1} . Consequently, Ω is equal to the union of all the subdomains Ω^l .

To build a hierarchical B-spline basis, a sequence of tensor-product B-spline spaces is introduced:

$$\mathcal{V}^0 \subset \mathcal{V}^1 \subset \mathcal{V}^2 \subset \mathcal{V}^3 \subset \dots \quad (9)$$

Each B-spline space \mathcal{V}^l has a corresponding basis \mathcal{B}^l and a corresponding mesh \mathcal{K}^l of tensor-product elements.

A hierarchical B-spline basis \mathcal{H} is constructed recursively based on the sequence of B-spline bases \mathcal{B}^l that span the domains Ω^l . In an initial step, the basis functions defined on the coarsest level, $l = 0$, are collected and assigned to \mathcal{H}^0 . The hierarchical B-spline basis \mathcal{H}^{l+1} is constructed by taking the union of all basis functions B in \mathcal{H}^l whose support is not fully enclosed

in Ω^{l+1} and all basis functions B in \mathcal{B}^{l+1} whose support lies in Ω^{l+1} . The recursive algorithm reads (Garau and Vázquez (2018)):

$$\left\{ \begin{array}{l} \mathcal{H}^0 =: \mathcal{B}^0 \\ \mathcal{H}^{l+1} =: \{B \in \mathcal{H}^l \mid \text{supp}(B) \not\subseteq \Omega^{l+1}\} \cup \\ \quad \{B \in \mathcal{B}^{l+1} \mid \text{supp}(B) \subseteq \Omega^{l+1}\}, \\ \quad \text{for } l = 0, \dots, n-2, \end{array} \right. \quad (10)$$

where the index l denotes the level of refinement. Basis functions collected in \mathcal{H} , where $\mathcal{H} := \mathcal{H}^{n-1}$, are called active, while basis functions in \mathcal{B}^l not present in \mathcal{H} are said to be inactive.

Associated with a hierarchical B-spline basis is a hierarchically refined mesh

$$\mathcal{K} := \bigcup_{l=0}^{n-1} \{K \in \mathcal{K}^l : K \in \Omega^l \text{ and } K \notin \Omega^{l+1}\} \quad (11)$$

wherein Ω^n is taken to be the empty set. An example of a hierarchically refined mesh associated with a two-dimensional hierarchical B-spline basis is displayed in Figure 6. A hierarchical B-spline basis is smooth over each element of its associated hierarchically refined mesh.

A hierarchical B-spline basis \mathcal{H} is illustrated for a one-dimensional example in Fig. 7. The top row shows a one-dimensional hierarchically refined mesh. Below the mesh, the basis functions for three refinement levels are shown. Following the recursion rule of Eq. (10), a B-spline basis \mathcal{H} is created through an initialization step with all bases in the subdomain Ω^0 refined to a level $l = 0$. All bases in the subdomain Ω^{l+1} with higher refinement level $l+1$ are added recursively, while existing basis functions of level l fully enclosed in Ω^{l+1} are discarded. The active B-spline basis functions \mathcal{H} are shown in black, while the inactive B-spline basis functions are shown in gray.

Hanging nodes are a byproduct of h -refinement in hierarchical refined meshes and are naturally handled by the B-spline bases. In contrast, hanging nodes in classic finite elements require extra treatments, such as the introduction of multi-point constraints.

3.4 Truncated B-splines

By construction, the hierarchical B-spline basis presented above does not fulfill the PU property. The truncated hierarchical B-spline basis constitutes an alternative to the hierarchical B-spline basis that does satisfy the PU property. Truncation also reduces the number of overlapping functions on adjacent hierarchical levels, see Giannelli et al. (2012). Considering a basis function

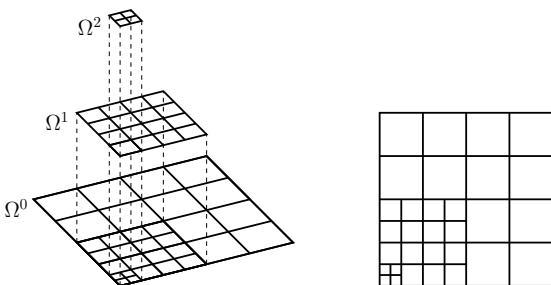


Fig. 6: Hierarchically refined mesh

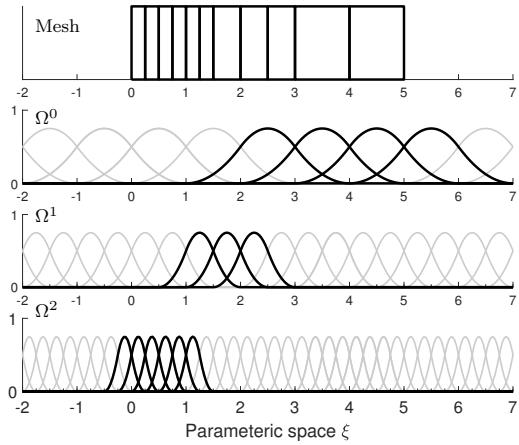


Fig. 7: One-dimensional hierarchical B-spline basis attained using three levels of refinement

B^l , part of \mathcal{B}^l and defined on the domain Ω^l , its representation in terms of the finer basis of level $l + 1$ is given as:

$$B^l = \sum_{B^{l+1} \in \mathcal{B}^{l+1}} c_{B^{l+1}}^{l+1} (B^l) B^{l+1}, \quad (12)$$

where $c_{B^{l+1}}^{l+1}$ is the coefficient associated to a basis function B^{l+1} .

The truncation of this basis function B^l , whose support overlaps with the support of finer basis functions B^{l+1} , part of \mathcal{B}^{l+1} and defined on Ω^{l+1} , is attained as follows (Giannelli et al. (2012); Garau and Vázquez (2018)):

$$\begin{aligned} \text{trunc}^{l+1}(B^l) &= \sum_{\substack{B^{l+1} \in \mathcal{B}^{l+1}, \\ \text{supp}(B^{l+1}) \not\subseteq \Omega^{l+1}}} c_{B^{l+1}}^{l+1} (B^l) B^{l+1}, \\ &= B^l - \sum_{\text{supp}(B^{l+1}) \subseteq \Omega^{l+1}} c_{B^{l+1}}^{l+1} (B^l) B^{l+1}. \end{aligned} \quad (13)$$

Following the creation of a hierarchical B-spline basis \mathcal{H} , a THB basis \mathcal{T} is constructed recursively by considering the truncation in Eq. (13), see Giannelli et al. (2012); Garau and Vázquez (2018):

$$\left\{ \begin{array}{l} \mathcal{T}^0 =: \mathcal{B}^0 \\ \mathcal{T}^{l+1} =: \{ \text{trunc}^{l+1}(B) \mid B \text{ in } \mathcal{T}^l \wedge \text{supp}(B) \not\subseteq \Omega^{l+1} \} \\ \quad \cup \{ B \in \mathcal{B}^{l+1} \mid \text{supp}(B) \subseteq \Omega^{l+1} \}, \\ \quad \text{for } l = 0, \dots, n-2. \end{array} \right. \quad (14)$$

The truncated basis \mathcal{T} spans the same space as the non-truncated basis \mathcal{H} and it admits a strong stability property (Giannelli et al. (2014)). Moreover, the smaller support results in a reduction in the number of nonzero basis functions per element and consequently a sparser system of linear equations in a finite element analysis.

The effect of the truncation is illustrated in Fig. 8. A univariate truncated and non-truncated basis \mathcal{T} and \mathcal{H} are juxtaposed. The first and second levels correspond to Ω^0 and Ω^1 respectively, while the bottom level represents the combination of the functions on these two levels. The comparison shows the reduced support of the truncated B-spline basis functions.

4 Implementation Details

This section details the implementation of the THB discretization described above. A computationally efficient approach is introduced to represent hierarchically refined meshes via a poly-tree data structure. Local mesh refinement strategies are outlined. A methodology for expressing THB basis functions associated with multiple, differently refined meshes in terms of Lagrange shape functions over each element in a union background is presented.

4.1 Poly-tree Data Structure for Hierarchically Refined B-spline Discretizations

The generation of a large hierarchical refined mesh can be expensive due to the high memory consumption, especially when creating all possible elements for each refinement level. The cost is compounded when creating multiple, differently refined hierarchical meshes. In this paper, we present a computationally and memory efficient approach to build differently refined meshes.

We first construct a Poly-Tree (PT) data structure, i.e., a quadtree in 2D and an octree in 3D. The PT data structure represents a set of hierarchically refined meshes. The tree depth corresponds to a refinement level l with the base level, 0, representing a coarse uniform tensor mesh. The nodes of this PT data structure are labeled PT cells and used to construct discretizations of state variable fields. Note that the PT cells are not associated with any specific set of basis functions.

PT cells with higher refinement levels $l > 0$ are created recursively. Starting from level 0, the PT cells are recursively subdivided into 4 and 8 PT cells in 2D and 3D, respectively. Considering a PT cell at a refinement level l , the PT cell at refinement level $l - 1$ from which the PT cell is created is referred to as its parent. The PT cells at refinement level $l + 1$ created by subdivision

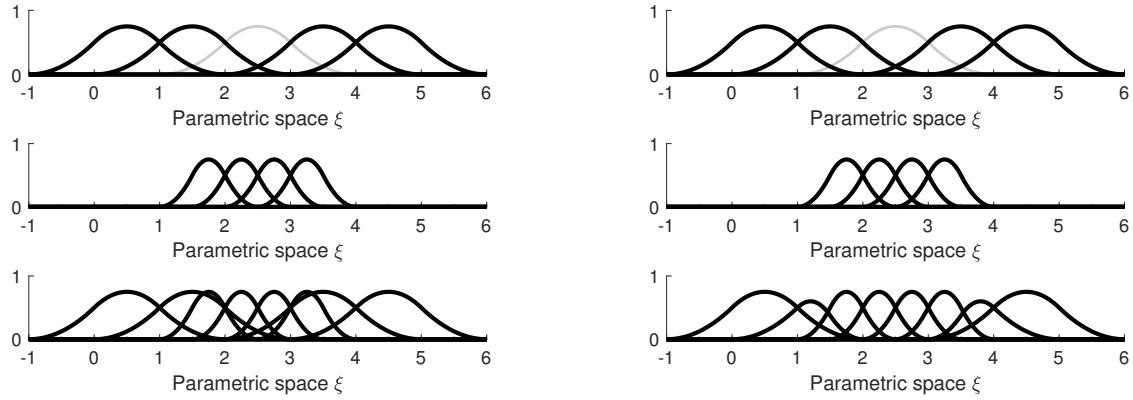


Fig. 8: Comparison of univariate HB-spline (left) and THB-spline basis functions (right)

of a PT cell at level l are called its children. The PT cells are only created once and only as needed.

To efficiently represent multiple, differently refined meshes with the same PT data structure, we introduce the concept of PT cell activation states. Each PT cell has multiple activation states which are represented by Activation Indices (AI). Possible activation states for a particular AI are either *active*, *refined*, or *inactive*. Active PT cells for a specific AI have *refined* parent PT cells and *inactive* children PT cells.

Although PT cells are only constructed once, the activation state concept allows for the creation of multiple, differently refined meshes. A single activation state can be stored as a binary number with only two bits. Thus, the activation states for all AIs of a PT cell are stored in $2 \cdot \# \text{AIs}$ bits. An illustration of a single quad-tree data structure with two AIs is presented in Fig. 9. The coloring of the PT cells indicates the activation state for a specific AI. For a particular AI, PT cells in blue, red, and white are *active*, *refined*, and *inactive*, respectively. Blue framed PT cells on the zero level illustrate the resulting mesh of active PT cells on a given AI. This enables the efficient construction of differently refined hierarchical B-spline discretizations based on a single PT data structure and a set of AIs. To see this, let Ω_I^l be the union of active and refined cells associated with level l and AI I . Then, we can create a hierarchical B-spline basis \mathcal{H}_I and a corresponding hierarchically refined mesh \mathcal{K}_I for AI I from the sequence

$$\Omega_I^{n-1} \subseteq \Omega_I^{n-2} \subseteq \dots \subseteq \Omega_I^0 = \Omega. \quad (15)$$

The resulting hierarchically refined mesh \mathcal{K}_I is precisely the set of active PT cells for AI I .

In our PT data structure, we only store PT cells that are either active or refined for at least one AI I . The element edge length $h_{\text{PTcell},m}^0$ for parametric direc-

tion m and on the coarsest level 0 is computed based on the given mesh size and the number of coarsest elements in each dimension, both of them are predefined by the user. The size of mesh elements corresponding to PT cells on levels $l > 0$ may be computed based on the coarsest element edge length and the refinement level via $h_{\text{PTcells},m}^l = 2^{-l} h_{\text{PTcell},m}^0$. Basis functions are assigned to mesh elements based on a pre-defined elemental basis function ordering. The elemental basis function ordering in this work follows the Exodus II standard (Shemon and Attaway (2014)).

4.2 Local Refinement Strategies

To generate a PT data structure that supports different hierarchically refined meshes, the PT is recursively refined for each AI and corresponding refinement criteria. The refinement procedure is described in Algorithm 1. For a specific AI, *active* PT cells are flagged for refinement based on chosen refinement criteria. Adjacent PT cells can also be flagged for refinement to increase the size of the refined region. Moreover, additional active PT cells for the current AI may be flagged considering mesh regularity requirements as discussed in Subsection 4.2.1. In case the PT data structure is generated in parallel using a domain decomposition strategy, flagged PT cells are communicated across adjacent subdomains to guarantee consistent refinement. Refinement is then performed by creating new children PT cells through subdividing all flagged PT cells unless the children PT cells do already exist. New PT cells are initialized with an *inactive* activation state for all AIs. For the AI currently considered for refinement, the activation state of the children PT cells is set to *active* while the parent PT cell is set to *refined*.

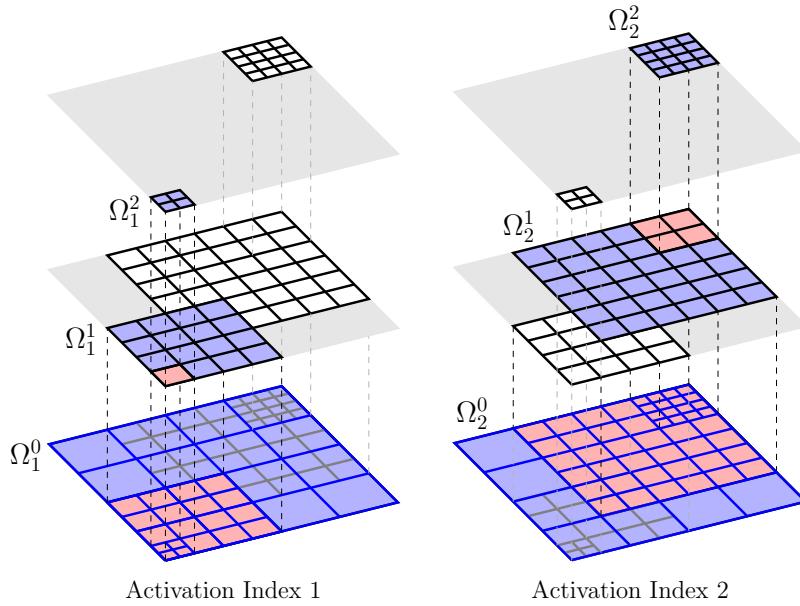


Fig. 9: One hierarchical refined quad-tree data structure with two sets of AIs

Algorithm 1 Refinement algorithm

```

1: for Activation Index (AIs) do
2:   Flag PT cells based on refinement criteria for AI
3:   Collect flagged PT cells into queue for refinement
4:   while changes in refinement queue do
5:     for all PT cells in queue for refinement do
6:       Apply refinement buffer Algorithm 2
7:       Collect flagged PT cells in queue for refinement
8:     end for
9:   end while
10:  Communicate refinement queue
11:  if Refined children PT cell does not exist then
12:    Create refined children PT cells
13:    Initialize new PT cells as inactive for all AIs
14:  end if
15:  Flag children PT cells as active for AI
16:  Flag parent PT cell as refined for AI
17: end for

```

4.2.1 Mesh Regularity Requirements

For construction of THB bases as discussed in Section 3, mesh regularity requirements need to be considered when constructing the PT data structure. The difference in refinement level between adjacent PT cells in the refined PT is limited to one. Furthermore, all active neighbor PT cells inside a so-called buffer zone of an active PT cell on level l must be of level greater or equal $l - 1$, see Eq. (14).

The buffer range $d_{\text{buffer},m}^l$ for a particular active PT cells is calculated by multiplying the PT cells size with a user-defined buffer parameter b_{buffer} , i.e., $d_{\text{buffer},m}^l = b_{\text{buffer}} h_{\text{PTcells},m}^l$. When creating a B-spline basis, the

width of the buffer zone in parametric direction m must be larger than or equal to the width of the basis function supports in parametric dimension m . In this work each state variable field is interpolated with an individual interpolation order p . To satisfy the mesh regularity requirement for all interpolation functions, the buffer parameter must be chosen as $b_{\text{buffer}} \geq p_{\max}$, where p_{\max} is the maximal polynomial degree of all used bases.

The refinement procedure for enforcing a buffer zone is summarized in Algorithm 2. The algorithm is applied to each PT cell initially flagged for refinement in Step 2 of Algorithm 1 and starts by determining its parent PT cell. The refinement status of the parent's neighbors, i.e., cells within the buffer range of the considered parent, are checked. If these neighbors are neither refined nor flagged for refinement, the distance d_m between the centers of the considered PT cell and its neighbor's children cells is calculated. This operation is trivial even if these neighboring children cells do not exist due to the PT data structure information. If any distance d_m is smaller than the buffer range $d_{\text{buffer},m}^l$, the particular neighbor cell is flagged for refinement. The algorithm is then applied recursively to all newly flagged neighbor cells, until no further cells are flagged for refinement. An efficient access to hierarchical mesh information, such as neighborhood relationships, is provided by the PT data structure discussed in Subsection 4.4. A visual representation of Algorithm 2 is provided in Fig. 10 for $b_{\text{buffer}} = 1$. PT cells with a green fill are cells flagged for refinement while PT cells with a green pattern fill are neighboring PT cells of the parent cell.

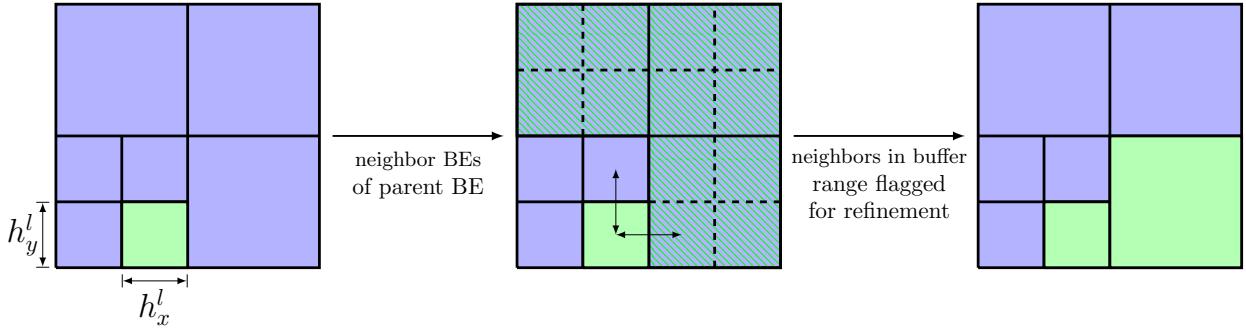


Fig. 10: Visualization of refinement buffer Algorithm 2

Algorithm 2 Refinement buffer algorithm

```

1: PT cell flagged for refinement in Algorithm 1
2: Get parent
3: Get parent's neighbors in buffer range  $d_{\text{buffer}}^l$ 
4: if neighbor is active and not flagged for refinement then
5:   Calculate distance  $d_{\text{ma}}$  considered cell and neighbors
   non-existing children PT cells
6:   if  $d_{\text{m}} < d_{\text{buffer}}$  then
7:     Flag neighbor for refinement
8:     Apply refinement buffer Algorithm 2 for neighbor
9:   end if
10: end if

```

4.3 Union Background Mesh and Extraction Operators

The framework presented above allows for a different hierarchically refined B-spline discretization for each state variable field. These discretizations are defined on locally refined background meshes that are not aligned with the external boundaries and material interfaces. To simplify the finite element formation and assembly process, a union¹ background mesh $\mathcal{K}_{\text{union}}$ is constructed from the hierarchically refined meshes for each state variable field as follows:

$$\mathcal{K}_{\text{union}} := \{K : K \in \mathcal{K}_I \text{ for some } I \text{ and there is no } K' \in \mathcal{K}_J \text{ for } J \neq I \text{ such that } K' \subsetneq K\}. \quad (16)$$

That is, an element of a hierarchically refined mesh for one of the state variable fields is an element of the union background mesh if it contains no finer element belonging to a hierarchically refined mesh for one of the other state variable fields. An example of a union background mesh is displayed in Fig. 3. The union background mesh is specially defined so that the hierarchical B-spline basis functions associated with each state variable field

¹ Note that the union background mesh is not attained via a set union of the separate hierarchically refined meshes. Instead, elements in this set union that fully contain smaller elements in the set union do not belong to the union background mesh.

are smooth over each element of the union background mesh. The union background mesh can be conveniently created via the PT data structure by combining all AIs used to generate THB background meshes. In particular, a PT cell corresponds to an element of the union background mesh if it is active for at least one activation index and active or inactive for all other activation indices. The union background mesh serves as the XFEM background mesh in which the geometry is immersed. Intersected XFEM background elements of the union background mesh are cut to generate a body-fitted integration mesh. Each THB discretization is enriched separately considering the B-spline interpolation and the immersed geometry represented on the union background mesh. For details on the enrichment strategy, the reader is referred to Noël et al. (2022). To increase the geometry resolution, the union mesh can be further refined, either locally or globally.

To facilitate operations performed on the union background mesh and on the integration mesh in the XFEM analysis, the THB basis functions for each locally refined background mesh are represented in terms of Lagrange shape functions on each element of the union background mesh via Lagrange extraction, see Schillinger et al. (2016); D'Angella et al. (2018). Namely, the THB basis functions for a locally refined background mesh can be represented element-wise in terms of Lagrange shape functions as

$$B_k(\xi) = \sum_j B_k(\xi_j) N_j^L(\xi) = \sum_j T_{jk}^L N_j^L(\xi) \quad (17)$$

where B_k is the k^{th} THB basis function, $\{N_j^L\}_j$ are the Lagrange shape functions over the element, $\{\xi_j\}_j$ are the locations at which the Lagrange shape functions $\{N_j^L\}_j$ are interpolatory, and $T_{jk}^L = B_k(\xi_j)$. We refer to T^L as a Lagrange extraction operator. The Lagrange shape functions over each element can in turn be represented in terms of Lagrange shape functions over a

child element as

$$N_j^L(\xi) = \sum_i N_j^L(\hat{\xi}_i) \hat{N}_i^L(\hat{\xi}_i) = \sum_i T_{ij}^h \hat{N}_i^L(\hat{\xi}_i) \quad (18)$$

where $\{\hat{N}_i^L\}_i$ are the Lagrange shape functions over the child element, $\hat{\xi}_i$ are the locations at which the Lagrange shape functions $\{\hat{N}_i^L\}_i$ are interpolatory, and $T_{ij}^h = N_j^L(\hat{\xi}_i)$. We refer to T^h as an h -refinement extraction operator. It follows that the THB basis functions can be expressed in terms of the Lagrange shape functions over the child element as

$$B_k(\xi) = \sum_i T_{ik} \hat{N}_i^L(\xi) \quad (19)$$

where

$$T_{ik} = \sum_j T_{jk}^L T_{ij}^h \quad (20)$$

is a Lagrange extraction operator that can be computed using the aforementioned Lagrange and h -refinement extraction operators. This process can be repeated to represent THB basis functions in terms of Lagrange shape functions and easily computable Lagrange extraction operators over descendent elements of the child element as well, thus enabling us to represent the THB basis functions for each state variable field in terms of Lagrange shape functions and Lagrange extraction operators over each element of the union background mesh. The extraction process from a quadratic B-spline basis to a quadratic, once refined Lagrange basis is illustrated in Fig. 11. Our framework limits the background element refinement to a factor of two, as presented in Subsection 4.2.1. This consequently limits the number of h -refinement extraction matrices to four in 2D and eight in 3D. These matrices can be precomputed and efficiently selected exploiting the PT data structure.

4.4 Background Mesh Data Structure

In this subsection, we discuss computational aspects of the PT data structure for storing and managing hierarchically refined meshes. This data structure aims at improving the overall computational efficiency and minimizes inter-process communication for parallel computations. The presented implementation is limited to quadtrees in 2D and octrees in 3D, which is sufficient for IFEMs.

Given the dimensions of the computational domain and the number of elements in each spatial direction, the base level 0 of the PT data structure is generated. Adopting a domain decomposition approach, the PT

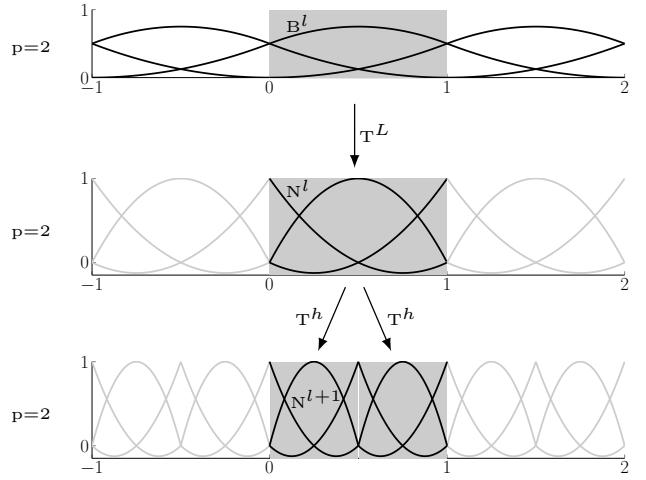


Fig. 11: Illustration of the extraction process

cells are grouped into subdomains. The mesh is hierarchically refined, and the union mesh is created in parallel on each processor. To facilitate the refinement process and the construction of the Lagrange extraction operators, we create overlapping subdomains to build an efficient PT data structure in parallel. In many cases a decomposition strategy assigning approximately the same number of PT cells to each subdomain is sufficient. However, as refinement can lead to a significant imbalance in number of PT cells across subdomains, it may be beneficial to choose a decomposition strategy that accounts for refinement.

At the coarsest level, 0, a unique PT cell ID can be calculated based on the global PT cell location as illustrated in the uppermost graphic of Fig. 13. With the ID of the parent PT cell at refinement level 0 and the position in the PT structure, unique IDs of refined PT cells can be determined directly, without inter-process communication and building local-to-global ID maps. The positions of the PT structure for a two dimensional twice refined PT cells are illustrated in Fig. 12.

Furthermore, the subdomain-local PT data structure allows for efficient access to cell neighborhood and hierarchy relationships. The PT data structure also speeds up the construction of extraction operators and may be used for the identification of basis functions. Basis functions can be uniquely identified by utilizing the PT structure in combination with the elemental Exodus II basis function index.

To reduce the communication across subdomains, an aura of PT cells is constructed around the set of PT cells owned by a specific processor, leading to overlapping subdomains. The width of the aura is chosen to be b_{buffer} PT cells on the coarsest refinement level with b_{buffer} defined as described in Subsection 4.2.1.

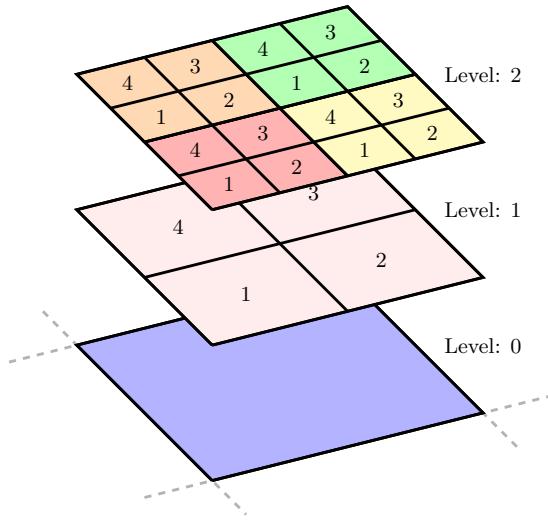


Fig. 12: A coarse two dimensional PT cell refined through recursive subdivision and resulting quadtree data structure

The PT cells in an aura are needed to define all THB basis functions that are nonzero over non-aura elements. Therefore, most computations can be performed on a subdomain level, interlaced with only a few inter-subdomain communications. This includes mesh refinement along the processor boundaries, the construction of Lagrange extraction operators, and the evaluation of face-oriented ghost stabilization residuals and Jacobians, see Section 5. To illustrate the parallel concept, Fig. 13 presents a global two dimensional domain of size 2×4 elements in blue. PT cells in red indicate aura cells. The global domain is decomposed into two processor local subdomains. Unique global PT cell IDs are calculated based on IDs of the coarsest refinement level. Red striped PT cells are aura cells shared with the neighboring processor. Blue striped PT cells are cells of the inverse aura. The green cells are flagged for refinement. Parallel consistent refinement is guaranteed through communication of the flagged PT cell IDs on aura and inverse aura.

5 Thermo-Elastic Model

The XIGA framework presented above is applicable to a broad range of physical problems that are described by partial differential equations. In this paper, we demonstrate the functionality of the developed framework with static elastic and thermo-elastic problems. The elastic model assumes infinitesimal strains and a linear elastic, isotropic material response. The thermal model accounts for linear diffusive heat transfer. The structural

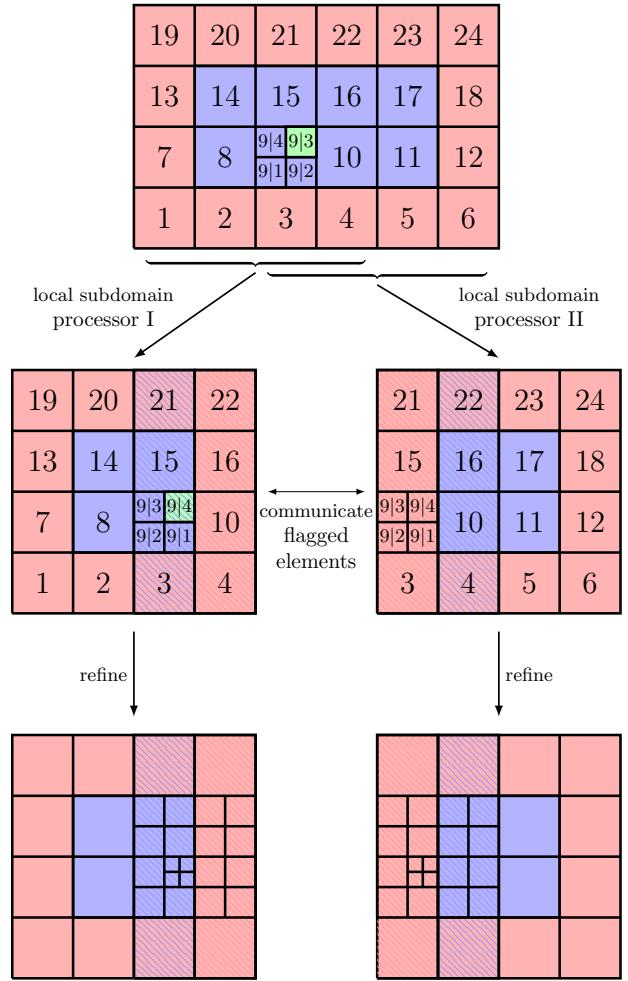


Fig. 13: Parallel local refinement of a global 2×4 domain

response depends on the temperature field through inelastic, isotropic thermal expansion. In this section, we present the variational form of the stabilized governing equations of the thermo-elastic model.

5.1 Governing Equations

The weak form of the residual is decomposed into the following four contributions:

$$\mathcal{R} = \mathcal{R}^T + \mathcal{R}^U + \mathcal{R}^N + \mathcal{R}^G = 0, \quad (21)$$

where \mathcal{R}^T and \mathcal{R}^U combine the volumetric contributions and contributions of Neumann boundary conditions to the thermal and structural residuals, respectively. The contributions of Nitsche's formulation for Dirichlet boundary conditions are collected in \mathcal{R}^N and the contributions of the face-oriented ghost stabilization in \mathcal{R}^G .

In the absence of body loads, the thermal and structural residual, \mathcal{R}^T and \mathcal{R}^U , are defined over all material domains Ω^m with $m = 1, \dots, N_m$ as follows:

$$\mathcal{R}^T = \sum_{m=1}^{N_m} \int_{\Omega^m} \nabla \delta \mathbf{T} \cdot \boldsymbol{\kappa} \nabla \mathbf{T} \, d\Omega - \int_{\Gamma_{N,T}^m} \delta \mathbf{T} \cdot \mathbf{q}_N \, d\Gamma, \quad (22)$$

and

$$\mathcal{R}^U = \sum_{m=1}^{N_m} \int_{\Omega^m} \delta \boldsymbol{\varepsilon}_u : \boldsymbol{\sigma} \, d\Omega - \int_{\Gamma_{N,u}^m} \delta \mathbf{u} \cdot \mathbf{f}_N \, d\Gamma, \quad (23)$$

where \mathbf{T} and $\delta \mathbf{T}$ are the temperature trial and test function, respectively. The isotropic thermal conductivity tensor is denoted by $\boldsymbol{\kappa}$. A heat flux \mathbf{q}_N is applied at the boundaries $\Gamma_{N,T}^m$. The displacement trial and test functions are denoted by \mathbf{u} and $\delta \mathbf{u}$, respectively. Traction forces, \mathbf{f}_N , are applied on the boundaries $\Gamma_{N,u}^m$. The Cauchy stress tensor is denoted by $\boldsymbol{\sigma} = \mathbf{D} \boldsymbol{\varepsilon}_e$ where \mathbf{D} is the isotropic constitutive tensor and $\boldsymbol{\varepsilon}_e$ is the total infinitesimal strain tensor with $\boldsymbol{\varepsilon}_e = \boldsymbol{\varepsilon}_u - \boldsymbol{\varepsilon}_T$. The mechanical strain $\boldsymbol{\varepsilon}_u$ is computed by $\boldsymbol{\varepsilon}_u = \frac{1}{2} (\nabla \mathbf{u} + \nabla \mathbf{u}^T)$. The thermal strain is denoted by $\boldsymbol{\varepsilon}_T = \alpha (\mathbf{T} - \mathbf{T}_0) \mathbf{I}$ where α is the CTE and T_0 is the reference temperature.

To weakly enforce Dirichlet boundary and interface conditions, we use an unsymmetric formulation of Nitsche's method. The associated residual, \mathcal{R}^N , is decomposed into boundary and interface terms, separately for the thermal and elastic subproblems as follows:

$$\mathcal{R}^N = \mathcal{R}_{\mathbf{T}}^{N,D} + \mathcal{R}_{\mathbf{u}}^{N,D} + \mathcal{R}_{\mathbf{T}}^{N,I} + \mathcal{R}_{\mathbf{u}}^{N,I}, \quad (24)$$

where the thermal and elastic contributions from Dirichlet boundary conditions are denoted by $\mathcal{R}_{\mathbf{T}}^{N,D}$ and $\mathcal{R}_{\mathbf{u}}^{N,D}$, respectively. The thermal and elastic contributions from interface conditions are $\mathcal{R}_{\mathbf{T}}^{N,I}$ and $\mathcal{R}_{\mathbf{u}}^{N,I}$, respectively. These Dirichlet boundary residual contributions are defined as follows:

$$\begin{aligned} \mathcal{R}_{\mathbf{T}}^{N,D} = & \sum_{m=1}^{N_m} - \int_{\Gamma_{D,T}^m} \delta \mathbf{T} \cdot (\boldsymbol{\kappa} \nabla \mathbf{T}) \cdot \mathbf{n}_\Gamma \, d\Gamma \\ & + \int_{\Gamma_{D,T}^m} (\boldsymbol{\kappa} \nabla \delta \mathbf{T}) \cdot \mathbf{n}_\Gamma (\mathbf{T} - \mathbf{T}_D) \, d\Gamma \\ & + \sum_{e=1}^{N_e^B} \gamma_{D,T}^e \int_{\Gamma_e^{D,T}} \delta \mathbf{T} \cdot (\mathbf{T} - \mathbf{T}_D) \, d\Gamma, \end{aligned} \quad (25)$$

and

$$\begin{aligned} \mathcal{R}_{\mathbf{u}}^{N,D} = & \sum_{m=1}^{N_m} - \int_{\Gamma_{D,u}^m} \delta \mathbf{u} \cdot \boldsymbol{\sigma}(\mathbf{u}) \mathbf{n}_\Gamma \, d\Gamma \\ & + \int_{\Gamma_{D,u}^m} \boldsymbol{\sigma}(\delta \mathbf{u}) \mathbf{n}_\Gamma \cdot (\mathbf{u} - \mathbf{u}_D) \, d\Gamma \\ & + \sum_{e=1}^{N_e^B} \gamma_{D,u}^e \int_{\Gamma_e^{D,u}} \delta \mathbf{u} \cdot (\mathbf{u} - \mathbf{u}_D) \, d\Gamma. \end{aligned} \quad (26)$$

The temperature T_D is prescribed on the boundary $\Gamma^{D,T}$ and the displacement \mathbf{u}_D is prescribed on $\Gamma^{D,u}$. The vector \mathbf{n}_Γ denotes the outward pointing normal on the boundary. The summations over all N_e^B elements in the union background mesh in Eqs. (25) and (26) penalize constraint violations along the boundaries $\Gamma_{D,T}^e$ and $\Gamma_{D,u}^e$ within the elements of the union background mesh. We henceforth refer to the elements of the union background mesh as background elements. The penalty parameters $\gamma_{D,T}^e$ and $\gamma_{D,u}^e$ depend on the size of the background element, h_e^B , and are defined as follows:

$$\gamma_{D,T} = c_{D,T} \frac{\kappa}{h_e^B} \quad \text{and} \quad \gamma_{D,u} = c_{D,u} \frac{E}{h_e^B}, \quad (27)$$

where κ is the isotropic material conductive and E is the Young's modulus of the linear elastic material. The parameters $c_{D,T} \geq 0$ and $c_{D,u} \geq 0$ control the accuracy of enforcing the Dirichlet boundary conditions.

Continuity of temperature and displacement fields and balance of heat flux and traction must be satisfied at all interfaces $\Gamma^{m,n} = \Omega^m \cap \Omega^n \neq \emptyset$. Nitsche's method for the thermal and structural interface conditions yields the following residual contributions:

$$\begin{aligned} \mathcal{R}_T^{N,I} = & \sum_{\Gamma^{m,n}} \sum_{e=1}^{N_e^B} - \int_{\Gamma_e^{m,n}} [\delta \mathbf{T}] \cdot \{\boldsymbol{\kappa} \nabla \mathbf{T}\} \cdot \mathbf{n}_{m,n} \, d\Gamma \\ & + \int_{\Gamma_e^{m,n}} \{\boldsymbol{\kappa} \nabla \delta \mathbf{T}\} \cdot \mathbf{n}_{m,n} [\mathbf{T}] \, d\Gamma \\ & + \gamma_{I,T}^e \int_{\Gamma_e^{m,n}} [\delta \mathbf{T}] [\mathbf{T}] \, d\Gamma, \end{aligned} \quad (28)$$

$$\begin{aligned} \mathcal{R}_u^{N,I} = & \sum_{\Gamma^{m,n}} \sum_{e=1}^{N_e^B} - \int_{\Gamma_e^{m,n}} [\delta \mathbf{u}] \cdot \{\boldsymbol{\sigma}(\mathbf{u})\} \mathbf{n}_{m,n} \, d\Gamma \\ & + \int_{\Gamma_e^{m,n}} \{\boldsymbol{\sigma}(\delta \mathbf{u})\} \mathbf{n}_{m,n} \cdot [\mathbf{u}] \, d\Gamma \\ & + \gamma_{I,u}^e \int_{\Gamma_e^{m,n}} [\delta \mathbf{u}] \cdot [\mathbf{u}] \, d\Gamma, \end{aligned} \quad (29)$$

where the jump operator is defined as $[\cdot] = (\cdot)^m - (\cdot)^n$. The numerical interface flux and traction are defined by the averaging operator as $\{\cdot\} = w^m(\cdot)^m + w^n(\cdot)^n$, where w^m and w^n are weights. The vector $\mathbf{n}_{m,n}$ denotes the normal vector on the interface pointing from phase m to phase n . The accuracy of enforcing the interface condition is controlled by the penalty terms with $\gamma_{I,T}^e$ and $\gamma_{I,u}^e$ being the elemental penalty factors.

We follow the work of Annavarapu et al. (2012) and define weights for the numerical heat flux as follows:

$$\begin{aligned} w_T^m &= \frac{\text{meas}(\Omega^m)/\kappa^m}{\text{meas}(\Omega^m)/\kappa^m + \text{meas}(\Omega^n)/\kappa^n}, \\ w_T^n &= \frac{\text{meas}(\Omega^n)/\kappa^n}{\text{meas}(\Omega^m)/\kappa^m + \text{meas}(\Omega^n)/\kappa^n}, \end{aligned} \quad (30)$$

and for the numerical traction as follows:

$$\begin{aligned} w_u^m &= \frac{\text{meas}(\Omega^m)/E^m}{\text{meas}(\Omega^m)/E^m + \text{meas}(\Omega^n)/E^n}, \\ w_u^n &= \frac{\text{meas}(\Omega^n)/E^n}{\text{meas}(\Omega^m)/E^m + \text{meas}(\Omega^n)/E^n}, \end{aligned} \quad (31)$$

where $\text{meas}(\Omega^k)$ is the surface in 2D or the volume in 3D of the domain occupied by the phase k within the background element. The elemental penalty factors in Eqs. (28) and (29) are defined by:

$$\begin{aligned} \gamma_{I,T}^e &= 2 c_{I,T} \frac{\text{meas}(\Gamma^{m,n})}{\text{meas}(\Omega^m)/\kappa^m + \text{meas}(\Omega^n)/\kappa^n} \\ \gamma_{I,u}^e &= 2 c_{I,u} \frac{\text{meas}(\Gamma^{m,n})}{\text{meas}(\Omega^m)/E^m + \text{meas}(\Omega^n)/E^n}, \end{aligned} \quad (32)$$

where the operator $\text{meas}(\Gamma^{m,n})$ measures the length in 2D or the area in 3D of the interface within the background element. The parameters $c_{I,T} \geq 0$ and $c_{I,u} \geq 0$ control the accuracy of enforcing the interface conditions.

5.2 Face-oriented ghost stabilization

The utilized immersed method may suffer from numerical instabilities caused by basis functions with small support within the geometric domain. This may occur when an interface moves close to the boundary of the support of a basis function. Such configurations may result in ill-conditioning of the system of linear equations as well as imprecise spatial gradients of the state variable field, see de Prenter et al. (2017). Recent works have proposed a variety of strategies to mitigate the issues arising from basis function with a small support. Most notable strategies are the basis function removal as presented by Elfverson et al. (2018), the concept of web-splines as presented by Höllig et al. (2001) and Marussig and Hughes (2018), or face-oriented ghost stabilization introduced by Burman and Hansbo (2014).

This work adopts the face-oriented ghost stabilization approach presented by Burman and Hansbo (2014) and adapted by Noël et al. (2022) to fit the basis function enrichment strategy described in Section 2.

Let Ω^m be a material subdomain phase m and K_{Ω^m} the set of background elements that have a non-empty intersection with Ω^m :

$$K_{\Omega^m} := \{K \in \mathcal{K}_{\text{union}} : K \cap \Omega^m \neq \emptyset\}. \quad (33)$$

We define $\mathcal{F}_{\text{int}}^m$ as the set of interior facets of K_{Ω^m} , i.e., the facets F shared between two background elements $\Omega_F^{m,+}$ and $\Omega_F^{m,-}$ of K_{Ω^m} . Let $\tilde{\Gamma}$ be the union of all material interfaces and geometric boundaries, here defined as material interfaces between void and non-void regions. The set of ghost facets for phase m is:

$$\mathcal{F}_{\text{ghost}}^m := \left\{ F \in \mathcal{F}_{\text{int}}^m : \Omega_F^{m,+} \cap \tilde{\Gamma} \neq \emptyset \text{ or } \Omega_F^{m,-} \cap \tilde{\Gamma} \neq \emptyset \right\}.$$

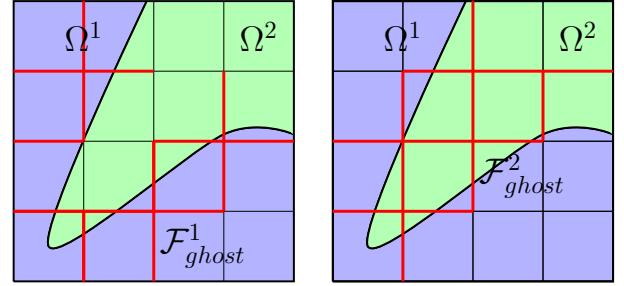


Fig. 14: Set of ghost facets $\mathcal{F}_{\text{ghost}}^1$ and $\mathcal{F}_{\text{ghost}}^2$ used for the face-oriented ghost stabilization for a two-material problem

$$(34)$$

Consider a ghost facet F^m shared between two adjacent background elements $\Omega_F^{m,+}$ and $\Omega_F^{m,-}$. The normal to the facet is \mathbf{n}_F and is chosen as $\mathbf{n}_F = \mathbf{n}_F^{m,+} = -\mathbf{n}_F^{m,-}$. The material layout subdivides the element $\Omega_F^{m,+}$ into $N_F^{m,+}$ connected subdomains $\Omega_{F,i}^{m,+}$ with $i = 1, \dots, N_F^{m,+}$ and the element $\Omega_F^{m,-}$ into $N_F^{m,-}$ connected subdomains $\Omega_{F,j}^{m,-}$ with $j = 1, \dots, N_F^{m,-}$.

We define $\mathbf{u}_{F,i}^{m,+}$ as the polynomial extension of the field $\mathbf{u}|_{\Omega_{F,i}^{m,+}}$ to all of \mathbb{R}^d and let $\mathbf{u}_{F,j}^{m,-}$ be the polynomial extension of the field $\mathbf{u}|_{\Omega_{F,j}^{m,-}}$ to all of \mathbb{R}^d .

Additionally, our formulation requires that: $|\partial\Omega_{F,i}^{m,+} \cap \partial\Omega_{F,j}^{m,-}| \neq 0$, and the ghost stabilization is only applied between $\mathbf{u}_{F,i}^{m,+}$ and $\mathbf{u}_{F,j}^{m,-}$ when the boundaries of $\Omega_{F,i}^{m,+}$ and $\Omega_{F,j}^{m,-}$, $\partial\Omega_{F,i}^{m,+}$ and $\partial\Omega_{F,j}^{m,-}$, respectively, meet along a portion of the facet F with a non-zero measure, e.g., the boundaries meet along more than a point in two dimensions and along more than a line in three dimensions.

With the above terminology defined, the contribution of the ghost stabilization for the displacement field to the residual equations is:

$$\mathcal{R}_{\mathbf{u}}^G = \sum_{m=1}^{N_m} \sum_{F \in \mathcal{F}_{\text{ghost}}^m} \sum_{i=1}^{N_F^{m,+}} \sum_{j \in \mathcal{J}_{F,i}^m} \left[\int_F \gamma_G^{\mathbf{u}} h^{\tilde{k}} [\![\partial_n^p \mathbf{u}]\!] [\![\partial_n^p \mathbf{u}]\!] d\Gamma \right], \quad (35)$$

where the set $\mathcal{J}_{F,i}^m$ is defined as:

$$\mathcal{J}_{F,i}^m := \{j \in \{1 \dots N_F^{m,-}\} : |\partial\Omega_{F,i}^{m,+} \cap \partial\Omega_{F,j}^{m,-}| \neq 0\}, \quad (36)$$

the jump operator $[\![\cdot]\!]$ is defined as:

$$[\![\partial_n^p \mathbf{u}]\!] = \left(\partial_n^p \mathbf{u}_{F,i}^{m,+} - \partial_n^p \mathbf{u}_{F,j}^{m,-} \right), \quad (37)$$

and

$$[\![\partial_n^p \delta \mathbf{u}]\!] = \left(\partial_n^p \delta \mathbf{u}_{F,i}^{m,+} - \partial_n^p \delta \mathbf{u}_{F,j}^{m,-} \right), \quad (38)$$

where $\partial_n^k(\cdot)$ is the k^{th} order normal derivative operator $\partial_n^k(\cdot) = \nabla^k(\cdot) \cdot \mathbf{n}_F$ where $\nabla^k(\cdot)$ is the k^{th} order spatial gradient. The parameter p is the polynomial order of the approximation, i.e., the THB discretization. It should be noted that only the p^{th} contribution is non-trivial for the THB discretization used in this work. The ghost penalty parameter is denoted γ_G^u and is defined as a multiple of the Young's modulus E of the considered material.

The ghost penalization for the temperature field is defined similarly as:

$$\mathcal{R}_T^G = \sum_{m=1}^{N_m} \sum_{F \in \mathcal{F}_{ghost}^m} \sum_{i=1}^{N_F^{m,+}} \sum_{\mathcal{J}_{F,i}^m} \left[\int_F \gamma_G^T h^{\tilde{k}} [\![\partial_n^p \mathbf{T}]\!] [\![\partial_n^p \mathbf{T}]\!] dI \right], \quad (39)$$

where the jump operator is defined as:

$$[\![\partial_n^p \delta \mathbf{T}]\!] = \left(\partial_n^p \delta \mathbf{T}_{F,i}^{m,+} - \partial_n^p \delta \mathbf{T}_{F,j}^{m,-} \right), \quad (40)$$

and

$$[\![\partial_n^p \mathbf{T}]\!] = \left(\partial_n^p \mathbf{T}_{F,i}^{m,+} - \partial_n^p \mathbf{T}_{F,j}^{m,-} \right). \quad (41)$$

The ghost penalty parameters, γ_G^u and γ_G^T , are defined as multiples of the Young's modulus E and the conductivity κ of the considered material, respectively. The parameter \tilde{k} is defined as $\tilde{k} = 2(p-1)+1$ and enables control over the influence of the ghost penalty term. Due to the application of C^{p-1} continuous THB bases, only jumps in gradients of order p must be penalized. Further details on how to choose the ghost penalty term can be found in Burman and Hansbo (2014). In this work we commonly choose a penalty parameter $\gamma_G = 0.001$.

6 Numerical Examples

In this section, we present 2D and 3D examples which illustrate the basic concepts and computational performance of the proposed discretization framework. Canonical examples show that our approach yields optimal convergence rates for linear, quadratic, and cubic B-spline approximations. An example with a stress singularity illustrates the computational advantages of locally refined discretizations. Finally, the analysis of a polycrystalline micro-structure demonstrates the applicability of our framework to complex multi-material problems.

To quantify the accuracy of the XFEM analyses, we define the error of a generic vector state variable field \mathbf{v} in the L^2 norm and H^1 semi-norm as:

$$\|\mathbf{v} - \mathbf{v}^h\|_{L^2} = \sqrt{\sum_{m=1}^{N_m} \int_{\Omega^m} |\mathbf{v} - \mathbf{v}^h|^2 d\Omega}, \quad (42)$$

and

$$|\mathbf{v} - \mathbf{v}^h|_{H^1} = \sqrt{\sum_{m=1}^{N_m} \int_{\Omega^m} |\nabla \mathbf{v} - \nabla \mathbf{v}^h|^2 d\Omega}, \quad (43)$$

where \mathbf{v}^h is the approximate field and \mathbf{v} is the reference solution which is either the analytical solution if available or a solution computed on a sufficiently refined discretization.

The systems of discretized governing equations are solved by the direct solver PARDISO for 2D problems, see Kourounis et al. (2018). A Generalized Minimal Residual (GMRES) method in combination with a dual threshold incomplete LU factorization with a degree of fill of 5.0 is used for 3D problems, see Saad (2003). The GMRES iterations are terminated if a relative drop of 10^{-10} of the normalized linear residual is achieved. All geometric and material parameters are given with each example in self-consistent units unless stated otherwise.

In the following subsections, we first present examples of single- and multi-material problems considering either a thermal or mechanical response. These examples characterize the fundamental features of the proposed analysis framework within single-physics settings. The last two examples consider coupled thermo-elastic single- and multi-material problems.

6.1 Two-Material Elastic Bar

This example studies the convergence rates of the proposed immersed B-spline discretization framework. We consider the 3D bar shown in Fig. 15. The bar has the dimension $1.0 \times 0.5 \times 0.5$. The left face of the bar is clamped, and the bar experiences a body load in axial direction $b_x = 2x^2$. The bar is composed of two linear elastic, isotropic materials, separated by an interface which is inclined against the x-axis. To facilitate comparison against an analytical solution, both materials are assigned the same properties: a Young's modulus $E = 1.0$ and a Poisson's ratio $\nu = 0.0$. Note that this setup allows for a 1D analytical model. The analytical displacement in x-direction is:

$$u(x) = \frac{L}{6EA} (4L^3x - x^4), \quad (44)$$

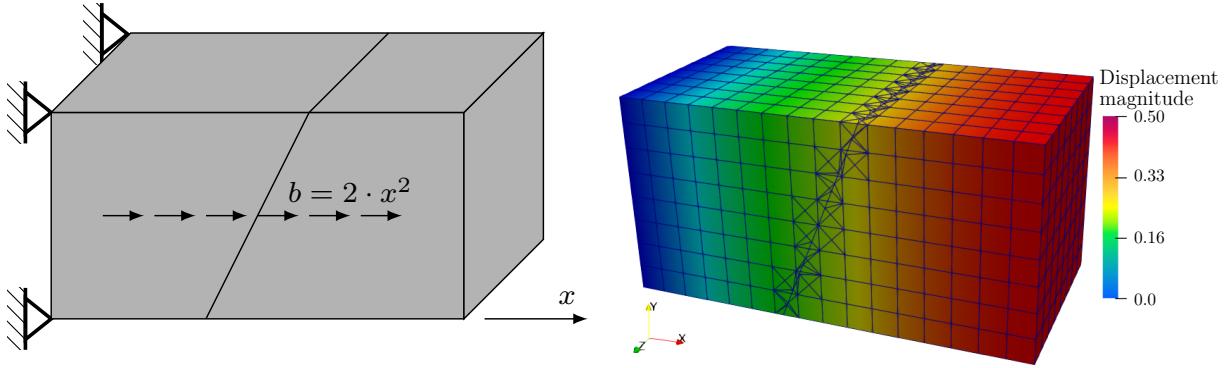


Fig. 15: Three dimensional bar problem: problem setup (left); contours of displacement magnitude evaluated on coarsest discretization (right)

where $L = 1.0$ is the bar length and $A = 0.25$ the cross-sectional area.

To study the influence of the intersection configuration on the finite element solution with B-spline background mesh refinement, we rotate the interface in ten steps between $\frac{\pi}{4}$ and $\frac{\pi}{2}$ degrees. For each orientation, we increase the number of background mesh elements from $8 \times 4 \times 4$ to $128 \times 64 \times 64$ B-spline elements through uniform mesh refinement. We repeat this study for linear, quadratic, and cubic B-spline discretizations.

In Fig. 16, we plot the mean error for all intersection configurations and for each approximation order over the B-spline element edge length as described in Eqs.(42) and (43). The error bars represent the standard deviation for all intersection configurations per order and mesh refinement. For each B-spline order, we visualize the convergence rate by the triangles inserted in Fig. 16. The numerical results show that the convergence rates of the L^2 error norm agree with the theoretical, optimal convergence rates of $p + 1$, where p is the polynomial B-spline order. Similarly, the theoretical, optimal convergence rate of p , as determined by Evans et al. (2009); Remacle et al. (2012), is achieved in the H^1 semi-norm for linear, quadratic, and cubic discretizations. Our results suggest that the proposed immersed finite element approach recovers optimal convergence rates for sufficiently smooth state variable fields with uniform mesh refinement.

6.2 Single-Material Thermal Diffusion Problem

Analysis problems often include regions where the state variable fields exhibit large spatial gradients. To reduce the local and global approximation errors, a fine discretization is needed in these regions. The following thermal diffusion example demonstrates the ability

of the proposed immersed finite element framework to construct and locally refine discretizations.

We solve a thermal diffusion problem as described in Eq. (22) in an L-shaped domain, assuming a single-material with an isotropic conductivity $\kappa = 1$, see Fig 17. For this problem, the exact solution in polar coordinates is given as (D’Angella et al. (2016)):

$$T(r, \theta) = r^{\frac{2}{3}} \sin\left(\frac{2}{3}\theta\right) \quad (45)$$

The exact solution is enforced weakly to the entire boundary of the physical domain. The origin of the coordinate system is located at the reentrant corner. Note that the spatial gradients increase for $r \rightarrow 0$ and are infinite at $r = 0$.

Using the proposed analysis framework, the physical domain is immersed into a rectangular computational domain as shown in Fig. 17. We compare the convergence rates for uniform h -refinement and for local h -refinement around the reentrant corner.

The coarsest THB background mesh has 6×6 elements. Sequences of THB discretizations are constructed by uniform or local refinement with up to 6 refinement levels, see Section 4. The offset Δ between the computational and physical domains is chosen as $\Delta = -0.127$. This offset guarantees that the L-shaped domain boundary does not align with the THB background mesh for any mesh refinement configuration.

For different B-spline order and mesh refinements, we plot the error in the L^2 norm versus the number of Degrees Of Freedom (DOFs) in Fig. 18. Less DOFs are required to meet a specific error requirement for local refinement when compared to uniform refinement. This suggests that local refinement may lead to a significant reduction in computational cost. Due to the singularity at the reentrant corner, optimal convergence rates with mesh refinement cannot be recovered. Moreover, this

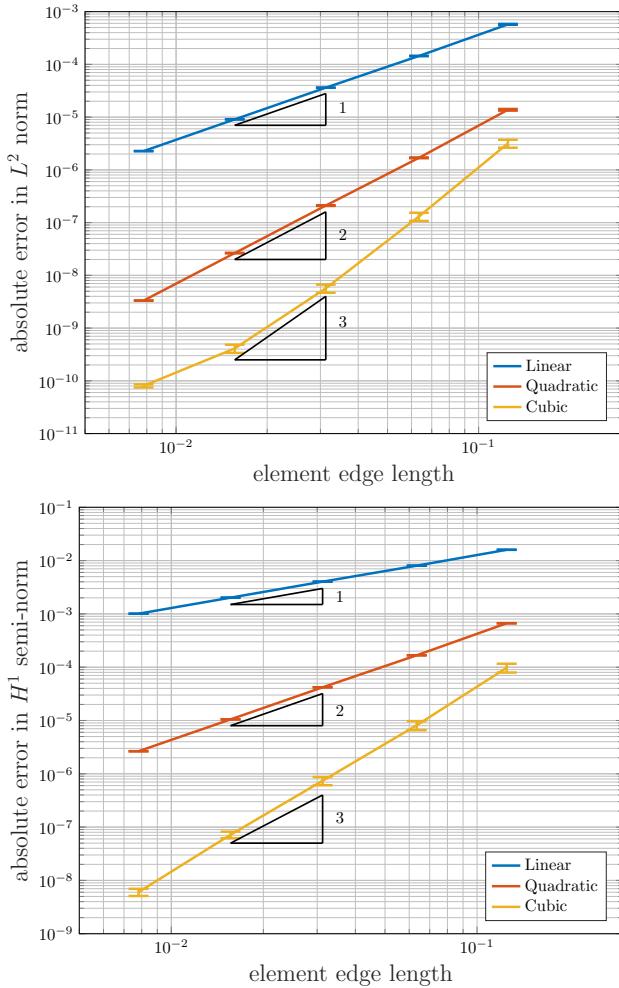


Fig. 16: Convergence rates in the L^2 error norm and H^1 semi-norm for the presented two-material bar problem

example demonstrates that the proposed framework allows for combining local h -refinement with higher order B-spline discretizations. For example, using cubic B-spline and local refinement leads to the lowest error for any number of DOFs. We do not present H^1 semi-norm error plots as they are not meaningful due to the singularity at the reentrant corner.

6.3 Polycrystalline Micro-structure

The proposed analysis framework is suited to model complex multi-material problems. While the previous examples established accuracy for single-material problems, this example demonstrates the multi-material capabilities of our approach and studies the computational cost associated with the generation of THB and union background meshes.

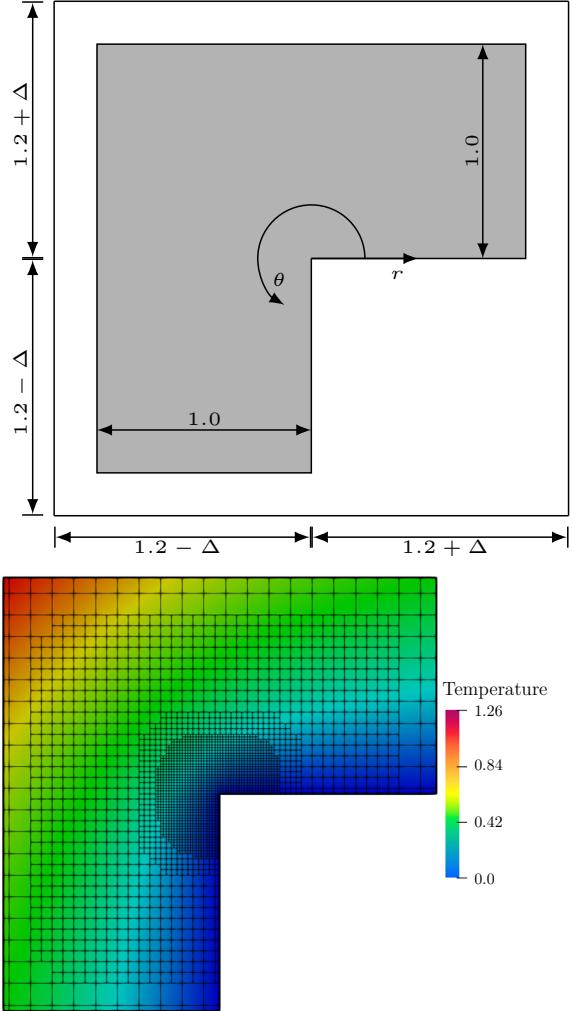


Fig. 17: Single-material thermal diffusion problem: L-shaped domain immersed into a rectangular computational domain (top); locally refined B-Spline discretization around the reentrant corner and temperature field (bottom)

We consider a representative volume element of a polycrystalline micro-structure and analyze its structural response. The edge length of the volume element is $150 \mu\text{m}$. The grain geometries as well as the grain material are defined through a 3D image file with $149 \times 149 \times 149$ voxels; see, for example, Rodgers (2015). The voxels define 471 individual grains. To demonstrate the multi-material capabilities, unique material properties are assigned to each grain. The Young's modulus E varies in the range of $[50 \cdot 10^3 \dots 500 \cdot 10^3] \text{ MPa}$ and the Poisson ratio ν in a range of $[0.25 \dots 0.35]$.

From the 3D image, we generate a level-set field for each phase such that the grain geometries are represented by the zero isocontours of the level-set fields. The grain geometries are immersed into a cubic domain, as

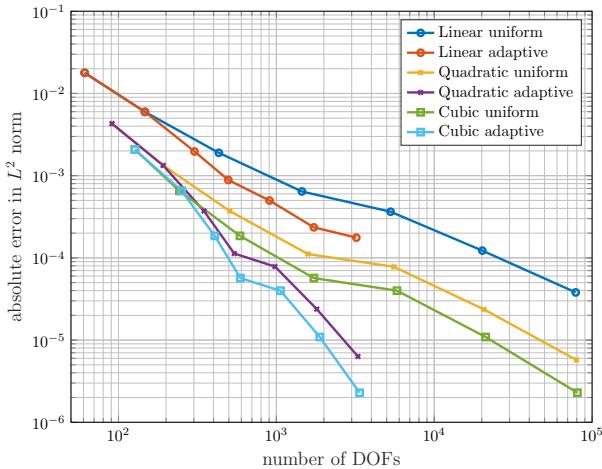


Fig. 18: Error convergence rate of temperature field with mesh refinement in the L^2 norm

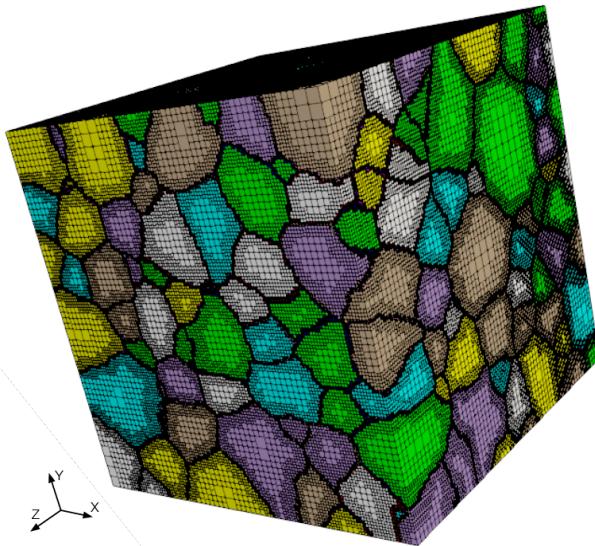


Fig. 19: Locally refined XFEM analysis model of a polycrystalline micro-structure

illustrated in Fig. 19. Note that the voxel-based grain shapes could be preprocessed to obtain smoother grain interfaces. However, this option is not utilized here to demonstrate the ability of the proposed analysis framework to operate directly on complex non-smooth voxel-based geometries.

The mechanical behavior of the polycrystalline micro-structure is modeled by linear elasticity with an isotropic material behavior. Perfect bonding is assumed at the grain interfaces. To demonstrate the ability of performing an XFEM analysis for this micro-structure, we apply a pressure load of $F = -4.0$ MPa on the $z = 0$ face and a zero-displacement Dirichlet boundary condition on the opposite face. The initial uniform B-spline back-

ground mesh has $32 \times 32 \times 32$ elements. In addition, two local refinement steps are performed at all grain boundaries. The displacement magnitude and Von Mises stress contours are shown in Fig. 20.

6.3.1 Scalability study

The PT data structure and domain decomposition strategy presented in Subsection 4.4 enable an efficient, parallel implementation of the proposed discretization framework. With the polycrystalline example, we study the overall performance and parallel scalability of this implementation.

For the scalability study, we start with $200 \times 200 \times 200$ PT cells at refinement level 0 and refine locally twice around the grain boundaries. For the demonstration of the scalability, we purposefully exploit a finer mesh than as in the previous analysis to guarantee a large ratio of interior to aura PT cells on each processor local subdomain for a large processor count. A small ratio of interior to aura elements negatively impacts the scalability due to increased inter-processor communication.

The resulting THB and union background meshes have each a total of 105,558,692 elements. The creation of the PT data structure and the derived discretizations can be subdivided into five distinct steps, see also Subsection 4. These steps include the refinement of the PT cells, the construction of the THB and union background meshes, the construction of the extraction operators, and the construction of facets needed for ghost stabilization.

To characterize the scalability of our implementation, we generate the THB and union background meshes in parallel, varying the number of subdomains from 4 to 160. Fig. 21 shows the execution time for the mesh generation only, i.e., the time needed for XFEM analyses is omitted. The computations are performed on four Intel Xeon Platinum 8160 "Skylake" nodes with 24 cores each, distributing the subdomains equally across all cores.

Mostly linear scaling with the processor count is observed for the refinement operation and the construction of the projection operators. Such behavior is expected as these steps do not need any inter-process communication. The construction of THB and union background meshes, as well as the construction of facets, show a linear scaling for a low processor count. They start to plateau with a high processor count, as these steps require communication of IDs in the aura. The aura size is based on refinement level l^0 and the buffer size as outlined in Section 4.4. Increasing the processor count while keeping the total domain size constant in-

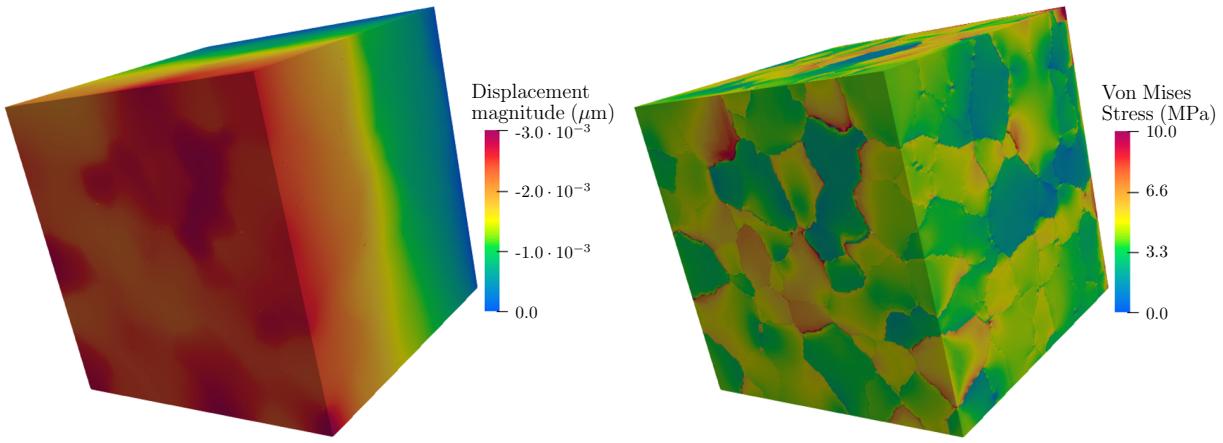


Fig. 20: XFEM analysis of polycrystalline micro-structure: Displacement magnitude contours (left); Von Mises stress contours (right)

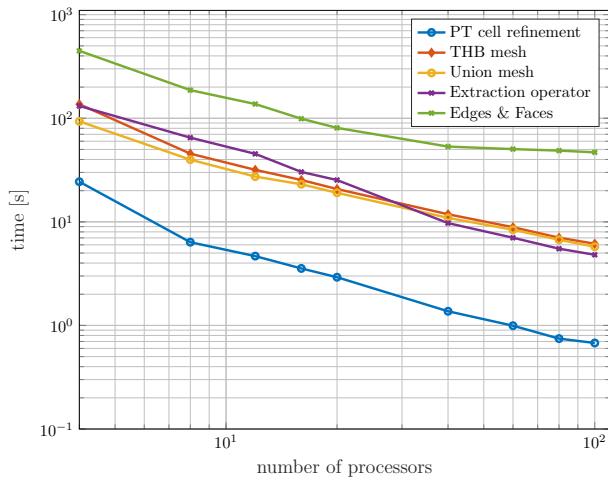


Fig. 21: Scalability study with processor count of the hierarchical meshing library, using $200 \times 200 \times 200$ PT cells at refinement level 0 and two refinement steps at grain boundaries

creases the ratio of aura PT cells to interior cells. This consequently increases communication and affects scalability.

The scalability study demonstrates that a mesh with over 10^8 elements can be generated on just four nodes with a total of 96 processors in less than 50 seconds. Furthermore, when creating a higher order discretizations, the computational time for the refinement of the background meshes and the construction of faces stays the same as neither are affected by the interpolation order. The construction of the discretizations, as well as the calculation of the extraction operators, needs slightly more memory and computational time as the support of higher order basis functions is increased.

6.4 Thermo-elastic Plate with Elliptic Hole

This example considers a single-material but multi-field configuration. We study the problem of a thermo-elastic plate with an elliptical inclusion under in-plane tension due to a thermal load as illustrated in Fig. 22. By exploiting symmetry, we only model a quarter of the domain. The setup consists of a solid, two-dimensional, square domain with a length of $L = 2.0$ and an elliptical inclusion at the origin with a semi-major axis of $A = 0.8136$ and a semi-minor axis of $B = 0.5753$. These values for the semi-axes are chosen such that the immersed interface does not align with the THB background mesh for all refinement levels. The temperature field is due to a uniform heat flux $q = 10.0$ applied along the elliptical hole and a prescribed temperature $T = 1.0$ at the right domain boundary. The material conductivity is $\kappa = 1.0$. The structural response is described by a linear thermo-elastic model and an isotropic constitutive behavior, with a Young's modulus $E = 1.0$, a Poisson's ratio $\nu = 0.3$, and a coefficient of thermal expansion $\alpha = 1.0$. The reference temperature is set to $T_0 = 0.0$.

To study the convergence rate of the XFEM solution with mesh refinement, we evaluate the thermal and structural response on a series of B-spline discretizations. Starting from the coarsest THB background mesh with 10×10 elements, uniformly refined B-spline discretizations with 20×20 , 40×40 , 80×80 , 160×160 and 320×320 elements are considered. The geometry of the elliptical inclusion is defined by an analytical level set function. A reference solution is computed numerically by calculating both fields on a uniformly refined B-spline mesh that is 7 times refined, 1280×1280 compared to the coarsest background mesh. Numerical studies lead by the authors have shown that the com-

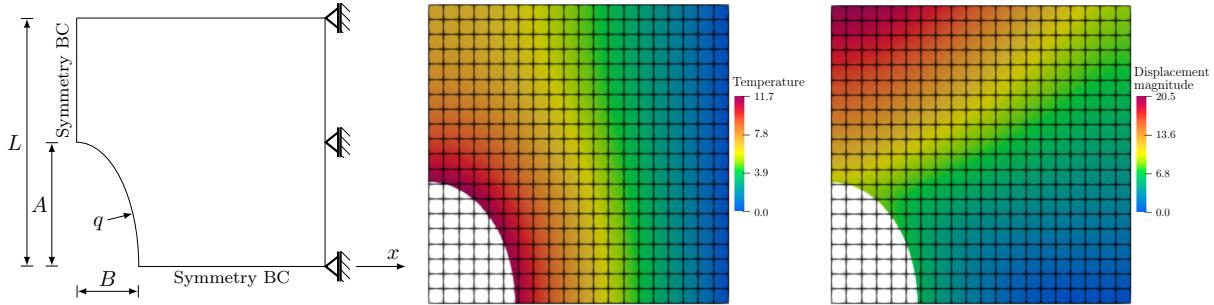


Fig. 22: Thermo-elastic plate with elliptic hole under thermal load: problem setup (left); temperature contours (center); displacement magnitude contours (right)

putation of the reference solution on a finer mesh does not noticeably improve the quality of the error computation.

For all B-spline discretizations, the union background mesh is 7 times uniformly refined relative to the coarsest background mesh. This eliminates the geometrical error between the coarse and the reference solution. Moreover, it simplifies the integration of the L^2 and H^1 errors as both the coarse and the reference solutions are represented on equally refined XFEM background meshes.

To illustrate the interdependence of discretization errors in multi-physics problems, we vary the refinement for the temperature discretization while the displacement field is always evaluated on the finest B-spline discretization, i.e., 320×320 elements. The contours of the absolute errors of the temperature and displacement fields are visualized in Fig. 23 for the case where the temperature field is discretized by linear B-splines on a 20×20 element mesh. This coarse temperature field discretization leads to errors in the interior of the elements. Therefore, the error visualization presents the characteristic tiled error patterns.

The L^2 and H^1 semi-norm errors of the temperature and the displacement fields are presented in Fig. 24. First, we consider the error of the temperature field. The convergence of the L^2 error norm for linear, quadratic, and cubic B-spline discretizations shows optimal convergence rate of $p + 1$. For the H^1 semi-norm, we observe optimal convergence rates of p . As expected, using a higher-order basis functions results in smaller errors for the temperature field.

Next, we consider the error of the displacement field, which is evaluated with linear, quadratic, and cubic basis functions and depends on the temperature field that is discretized by either linear, quadratic, and cubic B-splines. The errors of the displacement field in the L^2 norm and H^1 semi-norm are presented in Fig. 24. The convergence rate of the error with mesh refinement in

the L^2 norm is independent of the order of the displacement field basis functions and instead only depends on the order of the temperature field. Since the displacement field is discretized on a much finer mesh than the thermal field, the error of the displacement field is dominated by the error of the thermal field. The same observation can be made for the convergence rate in the H^1 semi-norm.

To gain further insight, we refine simultaneously the thermal and displacement fields with a specific difference in refinement level between both fields. The THB background element edge lengths of the thermal and displacement fields are denoted by H and h , respectively. The study is performed for a maximal difference in element size of $H = 4h$. This study is performed for linear, quadratic, and cubic interpolation orders for the thermal and the displacement fields. The setup of the thermal problem is identical the one presented above, see Fig. 24. The convergence rate of the displacement field with mesh refinement in the L^2 and H^1 semi-norm is presented in Fig. 25. We observe that the absolute error and convergence rate of the displacement problem only shows minimal differences when choosing an up to two times coarser thermal field than the displacement field.

Lastly, we present a convergence study where we allow for a difference in polynomial order. The thermal and displacement field are refined simultaneously such that both fields are on the same refinement level, i.e., $H = h$. Results are presented for a polynomial order of the displacement field p_U for $p_U = p_T$ and $p_U = p_T + 1$, where p_T is the order of the thermal field. The convergence rates with mesh refinement of the displacement field in the L^2 and H^1 semi-norm are presented in Fig. 26. We observe that choosing the thermal field one polynomial order lower than the displacement field results in the same convergence rate of the displacement field in the H^1 semi-norm. For this particular problem, theoretical results are not available for the rate of the

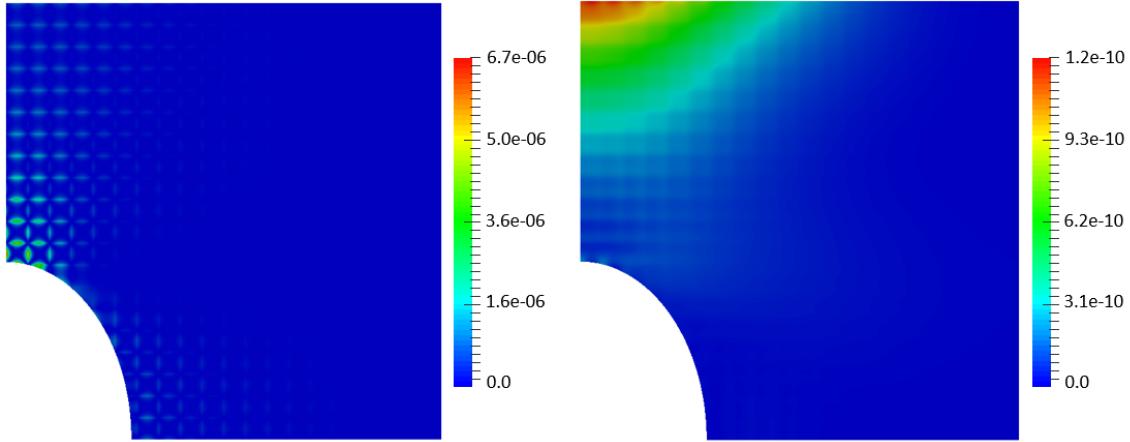


Fig. 23: Visualization of absolute error of temperature (left) and displacement field (right)

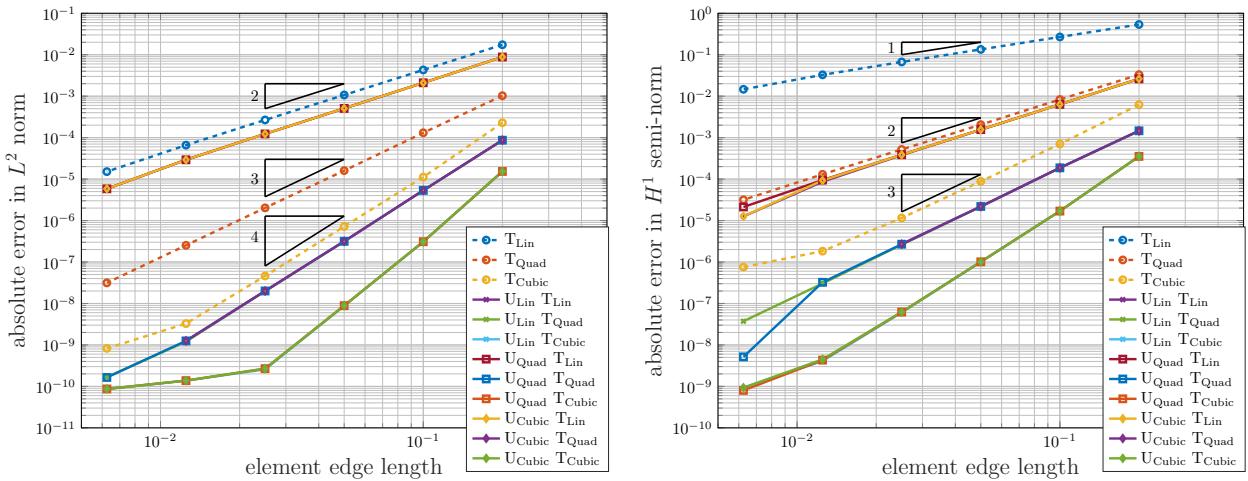


Fig. 24: Error convergence rate of temperature and displacement fields with mesh refinement in the L^2 norm (left) and H^1 semi-norm (right)

convergence for the displacement field in the L^2 norm, as the well-known Aubin-Nitsche technique cannot be applied in the thermo-elastic setting.

This example illustrates that for multi-physics problems, the discretization of individual fields needs to be performed in concert. The proposed discretization framework provides an efficient and flexible tool to select the interpolation order and mesh refinement level for individual fields to obtain a numerical solution that meets accuracy requirements with minimal computational costs, i.e., with minimal number of DOFs.

6.5 Thermo-elastic Multi-material Problem

This final example studies a multi-material, multi-physics configuration. For such problems, areas with large spatial gradients can vary based on the type of physics

and material. Differently locally refined discretizations for each field reduce the computational cost while simultaneously enabling for an accurate evaluation of the physical responses. To demonstrate this aspect, we consider the two-material, thermo-elastic problem of an expanding circle embedded in a non-expanding plate, as presented in the introduction in Fig. 1a. The circular inclusion is occupied by a material A and the plate by a material B. The temperature field is due to a spatially varying heat load of $q = 100.0 \cdot \sin(10.0 \cdot y) + 110.0$ applied along the left domain boundary and a prescribed temperature $T = 0.0$ at the right domain boundary. The conductivities of materials A and B are identical and are $\kappa_A = \kappa_B = 1.0$. The structural response of both materials is described by a linear thermo-elastic model and an isotropic, constitutive behavior, with Young's moduli $E_A = E_B = 1.0$, Poisson's ratios $\nu_A = \nu_B = 0.3$ and coefficients of thermal expansion $\alpha_A = 1.0 \cdot 10^{-5}$

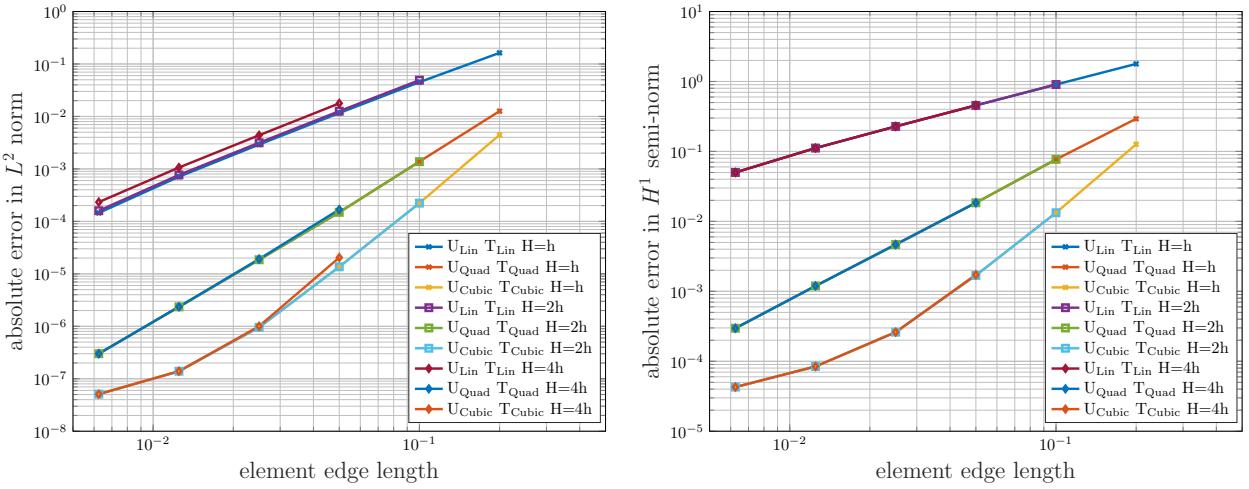


Fig. 25: Error convergence rate the displacement field with mesh refinement. L^2 norm (left) and H^1 semi-norm (right). Temperature and displacement field are refined simultaneously

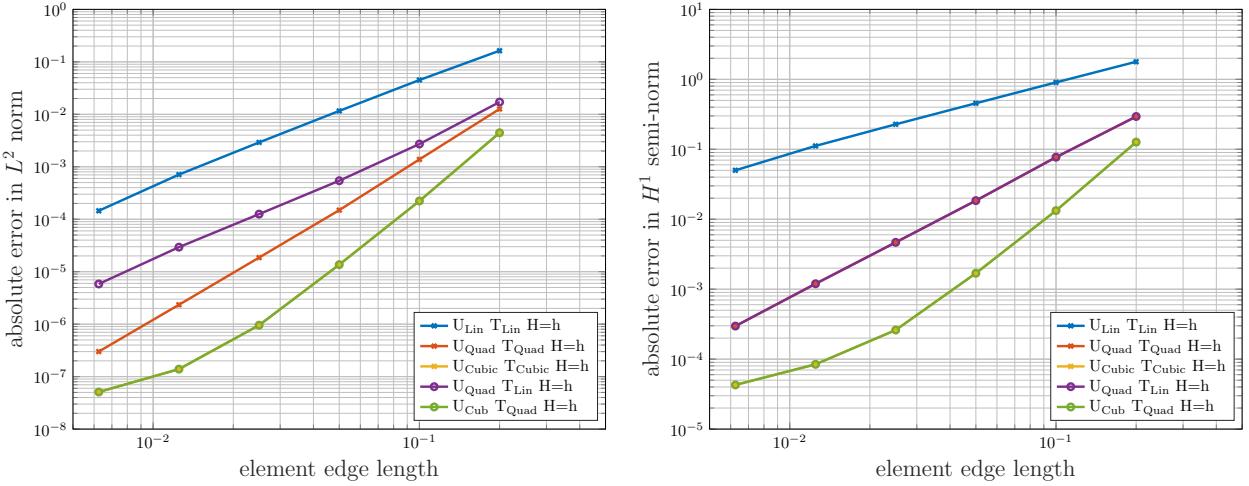


Fig. 26: Error convergence rate the displacement field with mesh refinement in the L^2 norm (left) and H^1 semi-norm (right) using identical refinements for the temperature and displacement fields

and $\alpha_B = 10.0 \cdot 10^{-5}$. The reference temperature is set to $T_0 = 0.0$. Perfect bonding is assumed at the interface between the two materials Γ^{AB} .

An XFEM analysis is performed to approximate the stress field at the material interface with high accuracy. We compare the temperature, displacement, Von Mises stress, and heat flux using uniformly and locally refined B-spline discretizations. A quadratic polynomial order is used for the temperature and displacement fields. The coarsest THB background mesh has 20×10 elements. For the uniform refinement case, the discretizations of the temperature and displacement fields are three times uniformly refined. In the local refinement case, the discretization for the temperature field is three times locally refined at the left domain boundary to accurately

represent the spatially varying load. In contrast, the discretization of the displacement field is three times locally refined around the circular material interface. The obtained heat flux and stress distributions are presented in Fig. 27 and show that the temperature and displacement fields exhibit large spatial gradients in these regions.

The number of DOFs resulting from the uniformly and locally refined THB background meshes are presented in Table 1. The locally refined discretization of the thermal field reduces the number of DOFs by a factor of ~ 17 when compared to the uniform discretization. When locally refining the displacement field, the number of DOFs is ~ 4.6 times smaller compared to the linear system associated with a uniform discretization.

Table 1: Comparison of the number of DOFs for temperature and displacement field for and locally uniformly refined discretizations

	Local Refinement	Uniform Refinement
#DOFs _{Temp}	818	13992
#DOFs _{Disp}	6020	27981

To accurately represent the geometry, the union background mesh is four times uniformly refined, irrespective of the B-spline discretizations of the temperature and displacement fields.

The displacement and temperature fields, as well as the Von Mises stress and the heat flux magnitude contours, are shown in Fig. 27. Qualitatively, the resulting fields for uniform and local refinement are equivalent.

The following four locally refined configurations are examined: (a) refinement for both the structural and thermal field at Γ^{AB} and at the Neumann boundary Γ^N ; (b) refinement of the structural field at Γ^{AB} and refinement of the thermal field at Γ^N ; (c) refinement of the structural field at Γ^{AB} ; and (d) refinement of the thermal field at Γ^N . In all cases the union background mesh is chosen to be uniformly refined.

Table 2 presents the L^2 error norm of the Von-Mises stress at the material interface Γ^{AB} for different local refinement configurations together with the size of the linear system. Local refinement for the first two Configurations, (a) and (b), yields the smallest error in the L^2 error norm. However, Configuration (b) results in a significant reduction of the computational cost as measured by the number of DOFs.

In addition, Configuration (c) yields a similar small error as Configuration (a) and (b). This behavior might be unexpected but can be explained with the fine uniform refinement of the union background mesh which allows for an accurate integration of the spatial varying load even with a coarse thermal discretization. Configuration (d) results in the highest error as only the thermal field is locally refined at the Neumann boundary.

7 Conclusion

This paper presents an immersed isogeometric finite element analysis framework with local mesh refinement based on a Heaviside enriched XFEM. Hierarchical, locally refined discretizations allow for refinement of the finite element approximations in regions of interest, balancing accuracy and computational cost. THB functions are utilized as they provide an elegant way to

Table 2: Comparison of error of the stress field at the interface Γ^{AB} in the L^2 error norm for different refinement configurations

Local Refinement	L^2 error norm	# DOFs
(a) Both fields at Γ^{AB} and Γ^N	1.2023e-6	10626
(b) Disp at Γ^{AB} and Temp at Γ^N	2.8881e-6	6838
(c) Disp at Γ^{AB}	2.8913e-6	5764
(d) Temp at Γ^N	3.8764e-4	2126

construct suitable, locally refined discretization spaces. Moreover, B-spline basis functions are an appealing choice over Lagrange basis functions because of their higher inter-element continuity and their increased computational efficiency. In multi-material, multi-physics problems, the resolution requirements may be different for individual state variable fields. The proposed framework allows for separate discretizations with different polynomial orders for each physical field. Furthermore, each discretization can be refined individually, both globally and locally, to meet field-specific accuracy requirements. In contrast to using the same polynomial order and refinement for all state variable fields, the proposed framework may lower significantly the computational cost. In this paper, THB background meshes are refined based on geometric refinement indicators. However, the framework permits any refinement indicators and can be used for adaptive mesh refinement strategies.

A PT data structure and mesh generation algorithms are presented for the efficient construction of differently refined meshes, both in terms of run time performance and memory needs. The concept of PT cell activation states enables using the same data structure to construct a set of different hierarchically refined THB background meshes. Using the PT data structure and the activation state concept, a union background mesh is constructed such that elements of the union background mesh are at the highest (or higher) refinement level of all corresponding elements of the THB background meshes. The union mesh supports the discretizations of all THB background meshes.

The union background mesh serves as the XFEM background mesh in which the geometry is immersed. In this paper, the geometry is represented by level set functions, and intersected elements are cut recursively by the zero isocontours of the level set fields. This process yields a single integration mesh which is aligned with the boundaries and interfaces defined by the level set functions. In this paper, the weak form of the governing equations is integrated by standard quadrature

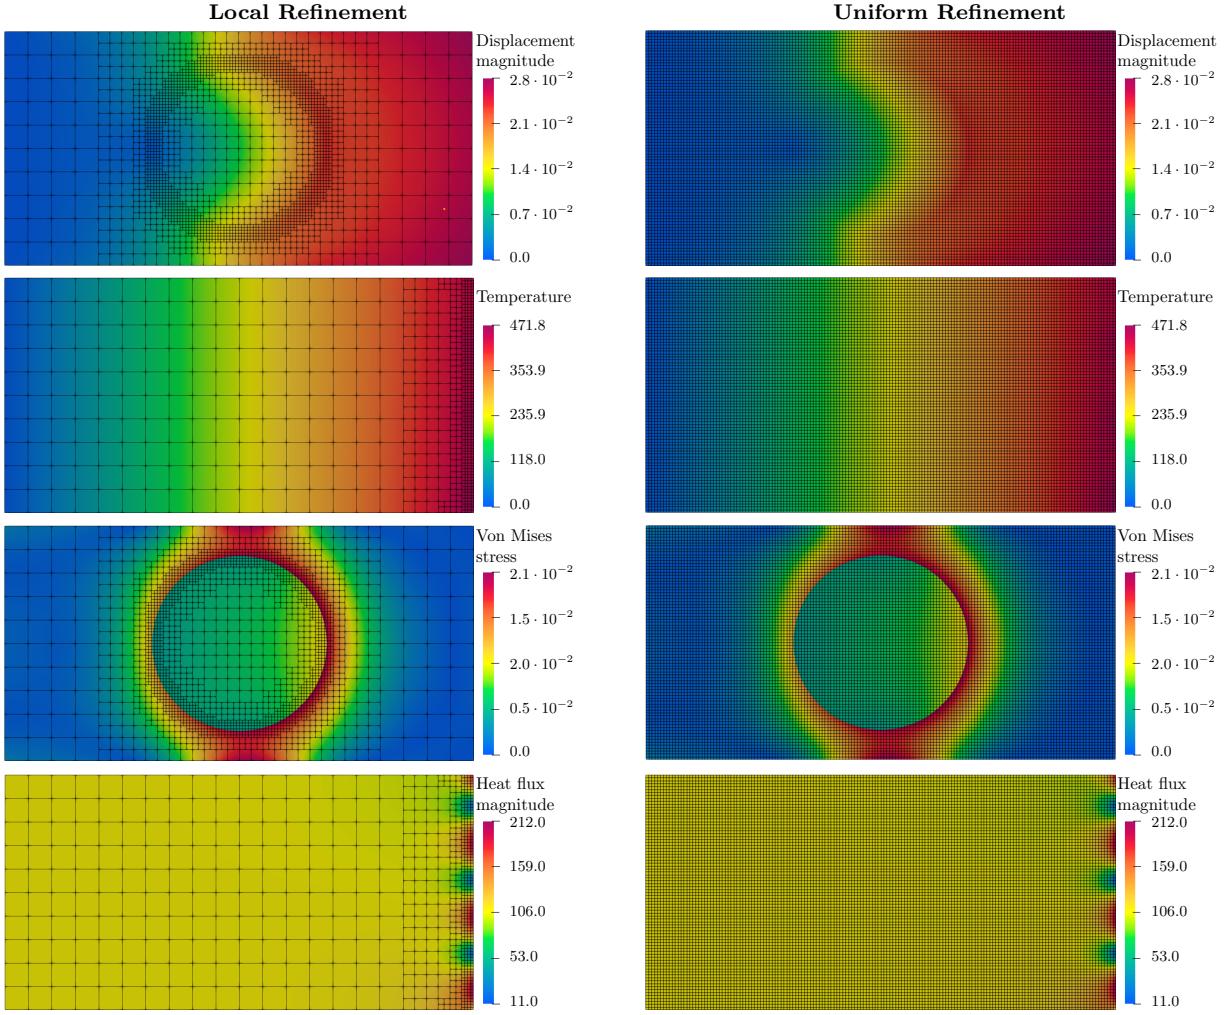


Fig. 27: Comparison of XFEM analysis results for locally and uniformly refined discretizations

rules on the integration mesh. However, other quadrature schemes can be applied, such as the ones proposed by Thiagarajan and Shapiro (2016) and Gundersen et al. (2021). The B-spline basis functions are enriched using the generalized Heaviside enrichment strategy of Noël et al. (2022). To facilitate operations performed on the union background mesh and on the integration mesh in the XFEM analysis, the THB basis functions are represented by Lagrange basis functions defined on the union background mesh via extraction operators. The union background mesh supports extraction operators for each THB discretization.

Numerical examples suggest that the proposed immersed isogeometric finite element framework generates discretizations that converge at theoretical, optimal convergence rates with mesh refinement. The application of our framework to a multi-material polycrystalline micro-structure shows that it is well suited to discretize complex multi-material problems in 3D. A

scalability study demonstrates that the proposed implementation scales with an increasing number of processors. The coupled thermo-elastic examples highlight the benefits of tailoring the discretization of individual state variable fields to their field-specific accuracy requirements.

Future work will focus on extending the proposed framework to utilize additional refinement criteria such as finite element error estimators. Moreover, more complex physics for which local mesh refinement is crucial, such as fluid flow at high Reynolds numbers, will be addressed.

Acknowledgment

The first, second, fourth, and fifth authors acknowledge the support for this work from the Defense Advanced Research Projects Agency (DARPA) under the

TRADES program (agreement HR0011-17-2-0022). The first author acknowledges partial auspice of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344 (LLNL-JRNL-842737). The third and fifth authors acknowledge the support of Sandia National Laboratories under PO 2120843. The fourth and fifth authors acknowledge the support of the National Science foundation under Grant 2104106. The opinions and conclusions presented in this paper are those of the authors and do not necessarily reflect the views of the sponsoring organization.

References

C. Annavarapu, M. Hautefeuille, and J. E. Dolbow. A robust nitsche's formulation for interface problems. *Computer Methods in Applied Mechanics and Engineering*, 225-228:44–54, 2012. ISSN 00457825. doi: 10.1016/j.cma.2012.03.008.

I. Babuška and J. M. Melenk. The partition of unity method. *International Journal for Numerical Methods in Engineering*, 40:727–758, 1997. ISSN 00295981. doi: 10.1002/(SICI)1097-0207(19970228)40:4(727::AID-NME86)3.0.CO;2-N.

Y. Bazilevs, V. M. Calo, J. A. Cottrell, J. A. Evans, T. J. Hughes, S. Lipton, M. A. Scott, and T. W. Sederberg. Isogeometric analysis using t-splines. *Computer Methods in Applied Mechanics and Engineering*, 199:229–263, 2010. ISSN 00457825. doi: 10.1016/j.cma.2009.02.036. URL <http://dx.doi.org/10.1016/j.cma.2009.02.036>.

T. Belytschko and T. Black. Elastic crack growth in finite elements with minimal remeshing. *International Journal for Numerical Methods in Engineering*, 45: 601–620, 1999.

T. Belytschko, R. Gracie, and G. Ventura. A review of extended/generalized finite element methods for material modeling. *Modelling and Simulation in Materials Science and Engineering*, 17, 2009. ISSN 09650393. doi: 10.1088/0965-0393/17/4/043001.

C. Bracco, A. Buffa, C. Giannelli, and R. Vázquez. Adaptive isogeometric methods with hierarchical splines: An overview. *Discrete and Continuous Dynamical Systems- Series A*, 39:241–262, 2019. ISSN 15535231. doi: 10.3934/dcds.2019010.

A. Buffa and C. Giannelli. Adaptive isogeometric methods with hierarchical splines: Optimality and convergence rates. *Mathematical Models and Methods in Applied Sciences*, 27:2781–2802, 2017. ISSN 02182025. doi: 10.1142/S0218202517500580.

E. Burman. La pénalisation fantôme. *Comptes Rendus Mathematique*, 348:1217–1220, 2010. ISSN 1631073X. doi: 10.1016/j.crma.2010.10.006. URL <http://dx.doi.org/10.1016/j.crma.2010.10.006>.

E. Burman. A penalty-free nonsymmetric nitsche-type method for the weak imposition of boundary conditions. *SIAM Journal on Numerical Analysis*, 50: 1959–1981, 1 2012. ISSN 0036-1429. doi: 10.1137/10081784X. URL <http://pubs.siam.org/doi/10.1137/10081784X>.

E. Burman and P. Hansbo. Fictitious domain methods using cut elements: Iii. a stabilized nitsche method for stokes' problem. *ESAIM: Mathematical Modelling and Numerical Analysis*, 48:859–874, 2014. ISSN 12903841. doi: 10.1051/m2an/2013123.

D'Angella, S. Kollmannsberger, E. Rank, and A. Reali. Multi-level bézier extraction for hierarchical local refinement of isogeometric analysis. *Computer Methods in Applied Mechanics and Engineering*, 328:147–174, 2018. ISSN 00457825. doi: 10.1016/j.cma.2017.08.017.

C. de Boor. On calculating with B-splines. *JOURNAL OF APPROXIMATION THEORY*, 6:50–62, 1972.

F. de Prenter, C. V. Verhoosel, G. J. van Zwieten, and E. H. van Brummelen. Condition number analysis and preconditioning of the finite cell method. *Computer Methods in Applied Mechanics and Engineering*, 316:297–327, 2017. ISSN 00457825. doi: 10.1016/j.cma.2016.07.006. URL <http://dx.doi.org/10.1016/j.cma.2016.07.006>.

S. C. Divi, C. V. Verhoosel, F. Auricchio, A. Reali, and E. H. van Brummelen. Error-estimate-based adaptive integration for immersed isogeometric analysis. *Computers and Mathematics with Applications*, 80:2481–2516, 2020. ISSN 08981221. doi: 10.1016/j.camwa.2020.03.026. URL <https://doi.org/10.1016/j.camwa.2020.03.026>.

C. A. Duarte, I. Babuši, and J. T. Oden. Generalized finite element methods for three-dimensional structural mechanics problems. *Computers and Structures*, 77:215–232, 2000.

A. Düster, J. Parvizian, Z. Yang, and E. Rank. The finite cell method for three-dimensional problems of solid mechanics. *Computer Methods in Applied Mechanics and Engineering*, 197:3768–3782, 2008. ISSN 00457825. doi: 10.1016/j.cma.2008.02.036.

D. D'Angella, N. Zander, S. Kollmannsberger, F. Frischmann, E. Rank, A. Schröder, and A. Reali. Multi-level hp-adaptivity and explicit error estimation. *Advanced Modeling and Simulation in Engineering Sciences*, 3, 12 2016. ISSN 22137467. doi: 10.1186/s40323-016-0085-5.

D. Elfverson, M. G. Larson, and K. Larsson. Cutiga with basis function removal. *Advanced Modeling and Simulation in Engineering Sciences*, 5:

1–19, 2018. ISSN 22137467. doi: 10.1186/s40323-018-0099-2. URL <https://doi.org/10.1186/s40323-018-0099-2>.

A. Embar, J. Dolbow, and I. Harari. Imposing Dirichlet boundary conditions with Nitsche's method and spline-based finite elements. *International Journal for Numerical Methods in Engineering*, 83:877–898, 8 2010. ISSN 00295981. doi: 10.1002/nme.2863. URL <http://doi.wiley.com/10.1002/nme.2863>.

J. A. Evans, Y. Bazilevs, I. Babuška, and T. J. Hughes. n-widths, sup-infs, and optimality ratios for the k-version of the isogeometric finite element method. *Computer Methods in Applied Mechanics and Engineering*, 198:1726–1741, 2009. ISSN 00457825. doi: 10.1016/j.cma.2009.01.021.

E. M. Garau and R. Vázquez. Algorithms for the implementation of adaptive isogeometric methods using hierarchical B-splines. *Applied Numerical Mathematics*, 123:58–87, 1 2018. ISSN 01689274. doi: 10.1016/j.apnum.2017.08.006.

A. Gerstenberger and W. A. Wall. An extended finite element method/Lagrange multiplier based approach for fluid-structure interaction. *Computer Methods in Applied Mechanics and Engineering*, 197:1699–1714, 2008. ISSN 00457825. doi: 10.1016/j.cma.2007.07.002.

C. Giannelli, B. Jüttler, and H. Speleers. THB-splines: The truncated basis for hierarchical splines. *Computer Aided Geometric Design*, 29:485–498, 2012. ISSN 01678396. doi: 10.1016/j.cagd.2012.03.025. URL <http://dx.doi.org/10.1016/j.cagd.2012.03.025>.

C. Giannelli, J. Bert, and H. Speleers. Strongly stable bases for adaptively refined multilevel spline spaces. *Adv Comput Math*, 40:459–490, 2014. doi: 10.1007/s10444-013-9315-2.

D. Gunderman, K. Weiss, and J. A. Evans. Spectral mesh-free quadrature for planar regions bounded by rational parametric curves. *Computer-Aided Design*, 130:102944, 2021. ISSN 0010-4485. doi: <https://doi.org/10.1016/j.cad.2020.102944>. URL <https://www.sciencedirect.com/science/article/pii/S0010448520301378>.

T. J. Hughes, J. A. Cottrell, and Y. Bazilevs. Isogeometric analysis: CAD, finite elements, NURBS, exact geometry and mesh refinement. *Computer Methods in Applied Mechanics and Engineering*, 194:4135–4195, 2005. ISSN 00457825. doi: 10.1016/j.cma.2004.10.008.

K. Höllig, U. Reif, and J. Wipper. Weighted extended b-spline approximation of dirichlet problems. *SIAM Journal on Numerical Analysis*, 39:442–462, 2001. doi: <https://doi.org/10.1137/s0036142900373208>.

D. Kourounis, A. Fuchs, and O. Schenk. Toward the next generation of multiperiod optimal power flow solvers. *IEEE Transactions on Power Systems*, 33:4005–4014, 2018. ISSN 08858950. doi: 10.1109/TPWRS.2017.2789187.

C. Lang, D. Makhija, A. Doostan, and K. Maute. A simple and efficient preconditioning scheme for heaviside enriched XFEM. *Computational Mechanics*, 54(5):1357–1374, 2014. ISSN 01787675. doi: 10.1007/s00466-014-1063-8.

D. Makhija and K. Maute. Numerical instabilities in level set topology optimization with the extended finite element method. *Struct Multidisc Optim*, 49:185–197, 2014. doi: 10.1007/s00158-013-0982-x.

B. Marussig and T. J. Hughes. *A Review of Trimming in Isogeometric Analysis: Challenges, Data Exchange and Simulation Aspects*, volume 25. Springer Netherlands, 2018. ISBN 1183101792209. doi: 10.1007/s11831-017-9220-9. URL <http://dx.doi.org/10.1007/s11831-017-9220-9>.

R. Mittal and G. Iaccarino. Immersed boundary methods. *Annual Review of Fluid Mechanics*, 37:239–261, 2005. ISSN 00664189. doi: 10.1146/annurev.fluid.37.061903.175743.

V. P. Nguyen. An introduction to standard and enriched isogeometric analysis through a simple matlab code isogeometric analysis : A review and computer implementation aspects. *Mathematics and Computers in Simulation*, 117:89–116, 2012. doi: 10.1016/j.matcom.2015.05.008.

J. Nitsche. Über ein variationsprinzip zur lösung von dirichlet-problemen bei verwendung von teilträumen, die keinen randbedingungen unterworfen sind. *Abhandlungen aus dem Mathematischen Seminar der Universität Hamburg*, 36:9–15, 1971. ISSN 18658784. doi: 10.1007/BF02995904.

L. Noël, M. Schmidt, C. Messe, J. A. Evans, and K. Maute. Adaptive level set topology optimization using hierarchical B-splines. *Structural and Multidisciplinary Optimization*, 62:1669–1699, 2020. ISSN 16151488. doi: 10.1007/s00158-020-02584-6.

L. Noël, M. Schmidt, K. Doble, J. A. Evans, and K. Maute. XIGA: An extended isogeometric analysis approach for multi-material problems. *Computational Mechanics*, 2022. ISSN 14320924. doi: 10.1007/s00466-022-02200-y.

J. Parvizian, A. Düster, and E. Rank. Finite cell method : h- and p-extension for embedded domain problems in solid mechanics. *Computational Mechanics*, 41:121–133, 2007. ISSN 01787675. doi: 10.1007/s00466-007-0173-y.

C. S. Peskin. The immersed boundary method. *Acta Numerica*, 11:479–517, 2002. ISSN 14740508. doi:

10.1017/S0962492902000077.

J. Remacle, J. Lambrechts, and B. Seny. Blossomquad: A non-uniform quadrilateral mesh generator using a minimum-cost perfect-matching algorithm. *International*, 104:1102–1119, 2012. doi: 10.1002/nme. URL <http://onlinelibrary.wiley.com/doi/10.1002/nme.3279/full>.

T. Rodgers. Exploration of process-structure linkages in simulated additive manufacturing microstructures, 2015. URL <https://doi.org/10.7910/DVN/KJMK9Z>.

Y. Saad. Iterative methods for sparse linear systems. *Iterative Methods for Sparse Linear Systems*, 2003. doi: 10.1137/1.9780898718003.

D. Schillinger and E. Rank. An unfitted hp-adaptive finite element method based on hierarchical b-splines for interface problems of complex geometry. *Computer Methods in Applied Mechanics and Engineering*, 200:3358–3380, 2011. ISSN 00457825. doi: 10.1016/j.cma.2011.08.002. URL <http://dx.doi.org/10.1016/j.cma.2011.08.002>.

D. Schillinger and M. Ruess. The finite cell method: A review in the context of higher-order structural analysis of cad and image-based geometric models. *Archives of Computational Methods in Engineering*, 22:391–455, 2015. ISSN 18861784. doi: 10.1007/s11831-014-9115-y. URL <http://dx.doi.org/10.1007/s11831-014-9115-y>.

D. Schillinger, L. Dedè, M. A. Scott, J. A. Evans, M. J. Borden, E. Rank, and T. J. R. Hughes. An isogeometric design-through-analysis methodology based on adaptive hierarchical refinement of NURBS, immersed boundary methods, and T-spline CAD surfaces. *Computer Methods in Applied Mechanics and Engineering*, 249-252:116–150, 2012. ISSN 0045-7825. doi: 10.1016/j.cma.2012.03.017. URL <http://dx.doi.org/10.1016/j.cma.2012.03.017>.

D. Schillinger, P. K. Ruthala, and Nguyen. Lagrange extraction and projection for nurbs basis functions: A direct link between isogeometric and standard nodal finite element formulations. *International*, pages 1102–1119, 2016. doi: 10.1002/nme. URL <http://onlinelibrary.wiley.com/doi/10.1002/nme.3279/full>.

E. R. Shemon and C. W. Attaway. User manual for exodus ii mesh converter, 2014. URL www.anl.gov.

T. Strouboulis, K. Copps, and I. Babuska. The generalized finite element method: an example of its implementation and illustration of its performance. *Int. J. Numer. Meth. Engng.*, 47:1401–1417, 2000. doi: 10.1002/(SICI)1097-0207.

K. Terada, M. Asai, and M. Yamagishi. Finite cover method for linear and non-linear analyses of heterogeneous solids. *International Journal for Numerical Methods in Engineering*, 58:1321–1346, 2003. ISSN 00295981. doi: 10.1002/nme.820.

V. Thiagarajan and V. Shapiro. Adaptively weighted numerical integration in the finite cell method. *Computer Methods in Applied Mechanics and Engineering*, 311:250–279, 2016.

C. V. Verhoosel, G. J. van Zwieten, B. van Rietbergen, and R. de Borst. Image-based goal-oriented adaptive isogeometric analysis with application to the micro-mechanical modeling of trabecular bone. *Computer Methods in Applied Mechanics and Engineering*, 284: 138–164, 2015. ISSN 00457825. doi: 10.1016/j.cma.2014.07.009. URL <http://dx.doi.org/10.1016/j.cma.2014.07.009>.

C. H. Villanueva and K. Maute. Density and level set-XFEM schemes for topology optimization of 3-D structures. *Comput Mech*, 54:133–150, 2014. doi: 10.1007/s00466-014-1027-z.