University of Glasgow

Elhaddad, M., Zander, N., Bog, T., Kudela, L., Kollmannsberger, S., Kirschke, J. S., Baum, T., Ruess, M. and Rank, E. (2018) Multi-level hp-finite cell method for embedded interface problems with application in biomechanics. *International Journal for Numerical Methods in Biomedical Engineering*, 34(4), e2951. (doi:10.1002/cnm.2951)

This is the author's final accepted version.

There may be differences between this version and the published version. You are advised to consult the publisher's version if you wish to cite from it.

http://eprints.gla.ac.uk/153359/

Deposited on: 11 December 2017

Enlighten – Research publications by members of the University of Glasgow
http://eprints.gla.ac.uk

Multi-level hp-finite cell method for embedded interface problems with application in biomechanics

Mohamed Elhaddad*[1], Nils Zander[1], Tino Bog[1], László Kudela[1], Stefan Kollmannsberger[1], Jan S. Kirschke[2], Thomas Baum[2], Martin Ruess[3] and Ernst Rank[1,4]

[1]Chair for Computation in Engineering, Technische Universität München, Arcisstr. 21, 80333 München, Germany
[2]Abteilung für Neuroradiologie, Klinikum rechts der Isar der Technischen Universität München, Ismaninger Str. 22, 81675 Munich, Germany
[3]School of Engineering, University of Glasgow, Rankine Building, Oakfield Ave. G12 8LT, United Kingdom
[4]Institute for Advanced Study, Technische Universität München, Lichtenbergstr. 2a, 85748 Garching, Germany

Abstract

This work presents a numerical discretization technique for solving three-dimensional material interface problems involving complex geometry without conforming mesh generation. The finite cell method (FCM), which is a high-order fictitious domain approach, is used for the numerical approximation of the solution without a boundary-conforming mesh. Weak discontinuities at material interfaces are resolved by using separate FCM meshes for each material sub-domain, and weakly enforcing the interface conditions between the different meshes. Additionally, a recently developed hierarchical hp-refinement scheme is employed to locally refine the FCM meshes in order to resolve singularities and local solution features at the interfaces. Thereby, higher convergence rates are achievable for non-smooth problems. A series of numerical experiments with two- and three-dimensional benchmark problems is presented, showing that the proposed hp-refinement scheme in conjunction with the weak enforcement of the interface conditions leads to a significant improvement of the convergence rates, even in the presence of singularities. Finally, the proposed technique is applied to simulate a vertebra-implant model. The application showcases the method's potential as an accurate simulation tool for biomechanical problems involving complex geometry, and it demonstrates its flexibility in dealing with different types of geometric description.

Keywords: high-order finite elements, hp-adaptivity, finite cell method, domain coupling, embedded interface problems, vertebra-implant model

Postprint submitted to *Int. J. Numer. Meth. Biomed. Engng.* November 21, 2017

*mohamed.elhaddad@tum.de, Corresponding author

Contents

1 Introduction

In engineering problems, an accurate resolution of domains with discontinuous material properties is often required to ensure a reliable design. Possible applications include the simulation of multi-phase materials, composite structures and biomechanical problems, where different tissues—often in combination with implants—are considered.

Problems of linear elasticity show a reduced regularity at the material interfaces, where the displacement field is in general only C^0-continuous. This fits perfectly to the classical finite element method (FEM) using a C^0-continuous approximation space [1, 2], as long as the material interface is matched exactly by edges (2D) or faces (3D) of the finite element mesh. However, this requires that the mesh generator is able to follow the material interfaces exactly, which can be cumbersome in cases of complex geometry. Moreover, high curvatures and kinks along the interfaces induce high solution gradients and singularities. This requires local mesh refinement for an accurate solution, which makes the task of mesh generation even more challenging.

Embedded or fictitious domain methods have emerged as an alternative to the boundary conforming FEM to avoid complex mesh generation. One example of this class of methods is the finite cell method (FCM) introduced in [3], which combines the embedded domain concept with high-order basis functions. The main idea is to extend the originally complex geometry by a fictitious domain to a much simpler shape, which can be easily meshed using a Cartesian grid. The use of high-order basis functions, which smoothly extend into the fictitious domain, allows the FCM to achieve exponential rates of convergence for smooth problems. The FCM has been investigated and successfully applied in various problems, including three-dimensional solid mechanics [4], linear thermoelasticity [5], geometrical non-linearities [6], implicit and explicit elastodynamics [7–10], and biomechanics [11, 12]. An open-source MATLAB® toolbox for the FCM has been presented in [13] offering an easy entry point into the topic.

Unfortunately, the standard FCM loses some of its attractive approximation properties for domains with multiple materials [14, 15]. There are two main challenges facing the FCM for embedded interface problems:

(1) The approximate solution within a finite cell is a polynomial (C^∞-continuous) whereas the exact solution is only C^0-continuous. The inability of the smooth polynomials to represent the discontinuity across the material interface yields an oscillatory solution and a reduction of the convergence rate.

(2) The use of a uniform Cartesian grid is not suitable for problems with locally high gradients or singularities, which are expected for material interfaces with complex geometry, as such problems require local mesh refinement.

To handle the first challenge, the numerical approximation needs to be able to reproduce a C^0-continuous solution. Several finite element approaches emerged that are suitable to handle weak discontinuities, *without* generating an interface-conforming mesh. Two major approaches are the partition of unity enrichment and the weak coupling of non-matching discretizations.

Partition of unity methods (PUM) [16], such as the generalized and extended finite element methods (GFEM, XFEM) [17, 18] enhance the numerical solution by enriching the basis with specially constructed functions to approximate the weakly discontinuous solutions [19, 20]. Recent developments in XFEM and GFEM for material interface problems include the use of adaptive mesh refinement in conjunction with the level-set method [21, 22], and the interface-enriched GFEM [23–27]. The local enrichment approach was combined with the FCM as presented by Joulaian and Düster for two-dimensional problems [15]. This combination leads to a significant improvement in the convergence rates of the FCM. However, the efficient and robust application of PUM enrichments to complex three-dimensional geometries remains a challenge.

The weak coupling approach allows for non-matching discretizations of the sub-domains, and it enforces the interface conditions in a weak sense. To that end, it is possible to apply the penalty method, Lagrange multipliers or Nitsche's method [28]. Nitsche's method is commonly used to weakly enforce interface conditions, as presented by Hansbo and Hansbo [29], and further developed by Dolbow and Harari [30–33]. The weak coupling approach has also been applied in

the context of the FCM, using Nitsche's method [34, 35], or based on a parameter-free coupling method [36]. There, it has been shown that a weak enforcement of the interface conditions enables the FCM to recovers its favorable convergence properties for piecewise smooth problems.

To address the second challenge, the finite cell mesh needs to be locally refined. Alternative versions of the FCM make use of unstructured meshes, and apply a local mesh refinement to resolve local solution features [37–39]. However, these mesh-based refinement approaches sacrifice the uniform grid structure of the FCM, which is particularly useful for image-based geometries that are common in biomechanical applications.

An attractive alternative is the application of the *hp*-version of the FEM for the non-uniform refinement of the underlying FCM meshes. The classical *hp*-FEM [40–42] is a powerful method which uses a combination of *h*- and *p*-refinement to efficiently approximate non-smooth solutions. The main idea is to sub-divide the elements close to the singularities in a recursive manner, in order to achieve a geometrical progression of element size. The resulting mesh has low order elements of small size close to the singularities, and high order elements of larger size further away from the singularities to approximate the smooth solution. While numerically efficient, *hp*-FEM for two- and three-dimensional problems requires to constrain the hanging nodes, which demands for a sophisticated discretization kernel [43, 44].

A simpler variant, which we employ in this work, is the recently introduced multi-level *hp*-version of the FEM [45, 46]. Instead of the classical refine-by-replacement concept, where an element is replaced by a set of smaller elements, the multi-level *hp*-FEM superposes the refined element with a hierarchy of overlay meshes. Using a simple formulation, in which homogeneous Dirichlet boundary conditions are enforced on the overlay meshes, the multi-level *hp*-FEM avoids the need to constrain hanging nodes. Thereby, it offers a simpler implementation for arbitrary hanging nodes, making it readily applicable to three-dimensional problems [47]. The multi-level *hp*-FEM has also been applied to cohesive fracture modeling [48], demonstrating its efficiency for propagating phenomena.

In this contribution, we apply the FCM in conjunction a weak enforcement of interface conditions and the multi-level *hp*-refinement scheme to solve material interface problems. This combination of these numerical techniques enables us to handle complex three-dimensional geometries, without the need for mesh generation. We demonstrate that the multi-level *hp*-scheme is very well-suited for the resolution of geometry induced stress concentrations, and presents a natural compliment to the FCM's strategy to avoid mesh generation. In particular, we demonstrate in this manuscript the use of the proposed refinement scheme in the context of biomechanical simulations, where a combination of heterogeneous tissue material properties and sharp material interfaces between tissue and osteosynthesis implants often requires a robust and flexible modeling approach to ensure reliable and accurate analysis results.

This paper is organized as follows: Section 2 offers a brief review of the finite cell method for embedded interface problems. In Section 3, we outline the multi-level *hp*-refinement scheme, and elaborate on the refinement strategy applied for material interface problems. In Section 4, we present a series of numerical experiments of embedded interface problems with geometry-induced stress concentrations. The results demonstrate the improved approximation accuracy of coupled FCM with multi-level *hp*-refinement. Finally, an application in biomechanics is presented in Section 5, where a thoracic vertebra with pedicle screws is simulated. The paper closes with a concluding outlook in Section 6.

2 The finite cell method for embedded interface problems

In this section, we present a brief review of the finite cell method for material interface problems, outlining the main concept of the method—including a coupling formulation that allows to weakly impose the interface conditions among sub-domains.

2.1 Model problem

We consider an open and bounded domain, denoted as the physical domain $\Omega_{\text{phy}} \subset \mathbb{R}^d, d \in \{2, 3\}$, with the boundary $\partial\Omega_{\text{phy}} = \Gamma_D \cup \Gamma_N$, and $\Gamma_D \cap \Gamma_N = \emptyset$ [34]. The physical domain Ω_{phy} is subdi-

vided into n disjoint, open, and bounded subsets $\Omega^{(k)}$

$$\Omega_{\text{phy}} = \bigcup_{k=1}^{n} \Omega^{(k)}, \quad \Omega^{(i)} \cap \Omega^{(j)} = \emptyset \text{ for } i \neq j. \tag{1}$$

The sub-domains $\Omega^{(k)}$ are separated by internal boundaries Γ_{ij}, denoting the material interfaces between $\Omega^{(i)}$ and $\Omega^{(j)}$. The model is governed by the partial differential equations of linear elastostatics, given in the strong form of the boundary value problem by:

$$\nabla \cdot \boldsymbol{\sigma}^{(i)} + \boldsymbol{b}^{(i)} = \mathbf{0} \qquad \forall \boldsymbol{x} \in \Omega^{(i)}, \quad i = 1, \dots, n \tag{2a}$$

$$\boldsymbol{\sigma}^{(i)} = \mathbb{C}^{(i)} : \boldsymbol{\varepsilon}(\boldsymbol{u}^{(i)}) \qquad \forall \boldsymbol{x} \in \Omega^{(i)}, \quad i = 1, \dots, n \tag{2b}$$

$$\boldsymbol{\varepsilon}(\boldsymbol{u}^{(i)}) = \frac{1}{2}\left(\nabla \boldsymbol{u}^{(i)} + \nabla \boldsymbol{u}^{(i)^\top}\right) \qquad \forall \boldsymbol{x} \in \Omega^{(i)}, \quad i = 1, \dots, n, \tag{2c}$$

where $\boldsymbol{\sigma}^{(i)}$ is the stress tensor, $\boldsymbol{b}^{(i)}$ is the body load, $\mathbb{C}^{(i)}$ is the elastic material tensor, $\boldsymbol{\varepsilon}$ is the strain tensor, and $\boldsymbol{u}^{(i)}$ is the displacement vector of the sub-domain $\Omega^{(i)}$. The system is subject to the boundary conditions:

$$\boldsymbol{u}^{(i)} = \hat{\boldsymbol{u}}^{(i)} \qquad \forall \boldsymbol{x} \in \Gamma_D^{(i)}, \quad i = 1, \dots, n \tag{3a}$$

$$\boldsymbol{\sigma}^{(i)} \cdot \boldsymbol{n}^{(i)} = \hat{\boldsymbol{t}}^{(i)} \qquad \forall \boldsymbol{x} \in \Gamma_N^{(i)}, \quad i = 1, \dots, n, \tag{3b}$$

where $\hat{\boldsymbol{u}}^{(i)}$ are the prescribed displacements on the Dirichlet boundary, $\hat{\boldsymbol{t}}^{(i)}$ are the prescribed tractions on the Neumann boundary, and $\boldsymbol{n}^{(i)}$ denotes the outward facing normal vector on the sub-domain's boundary $\partial\Omega^{(i)}$ (see Figure 1), and

$$\Gamma_D^{(i)} = \Gamma_D \cap \partial\Omega^{(i)} \tag{4a}$$

$$\Gamma_N^{(i)} = \Gamma_N \cap \partial\Omega^{(i)}. \tag{4b}$$

Additionally, the following interface conditions apply along Γ_{ij}:

$$\boldsymbol{u}^{(i)} = \boldsymbol{u}^{(j)} \qquad \forall \boldsymbol{x} \in \Gamma_{ij}, \tag{5a}$$

$$\boldsymbol{\sigma}^{(i)} \cdot \boldsymbol{n}^{(i)} = -\boldsymbol{\sigma}^{(j)} \cdot \boldsymbol{n}^{(j)} \qquad \forall \boldsymbol{x} \in \Gamma_{ij}. \tag{5b}$$

The first interface condition ensures the kinematic compatibility between the sub-domains, whereas the second condition prescribes the equilibrium of normal traction across the interfaces.

2.2 Aspects of the finite cell method

The finite cell method is motivated by the superior convergence properties achievable by the p-version of FEM. The geometry-conforming p-FEM requires special attention to mesh generation, which can become prohibitive for problems involving complex geometries. The FCM avoids the need for complex mesh generation procedures by combining the p-FEM with an embedded domain approach.

The principal idea of the FCM is to augment the physical domain Ω_{phy} by a fictitious domain Ω_{fict} to form a simple-shaped domain $\hat{\Omega} = \Omega_{\text{phy}} \cup \Omega_{\text{fict}}$. In order to recover the original problem, the influence of the fictitious domain is attenuated by penalization of the material parameters in the fictitious domain:

$$\mathbb{C}_{\text{fict}} = \epsilon \cdot \mathbb{C}_{\text{phy}}, \quad \epsilon \ll 1. \tag{6}$$

The modified problem is solved over the extended domain, which can be easily meshed with a Cartesian grid of finite cells. The FCM makes use of high-order basis functions which smoothly extend over the embedding domain to approximate the physical solution field. Typical choices of basis functions are the integrated Legendre polynomials, commonly used in the p-FEM [3],

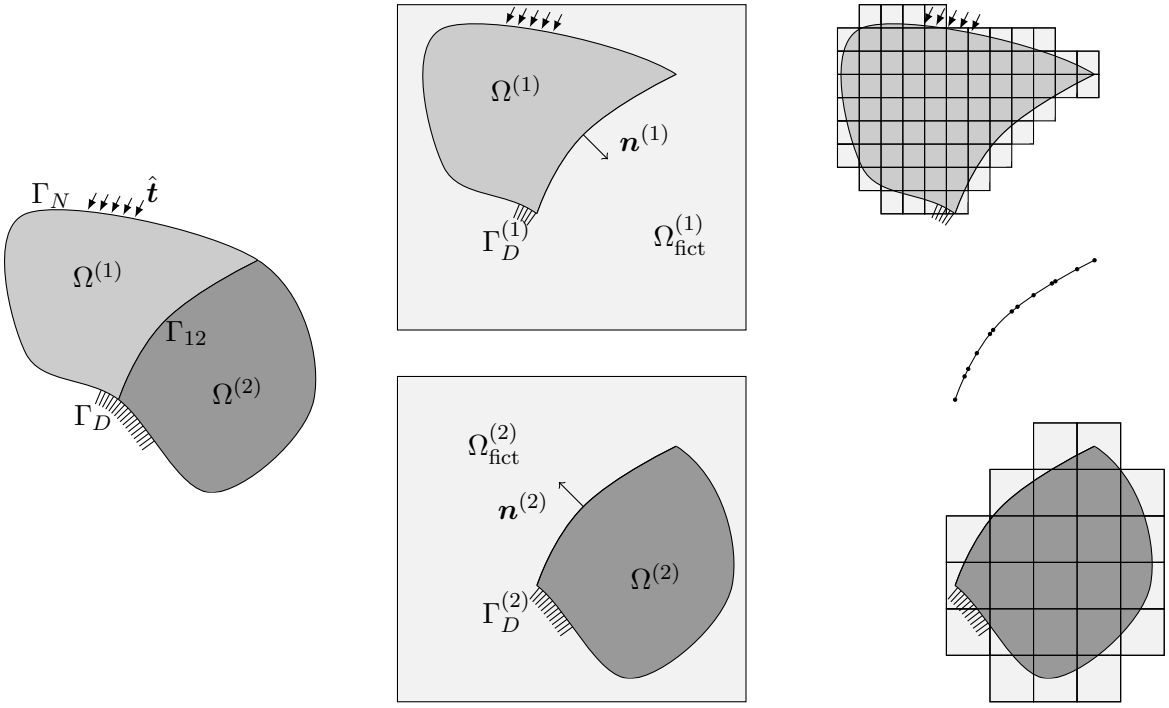

Figure 1: Coupling of non-matching FCM meshes, following [34].

and B-Splines used in Isogeometric analysis [6]. For applications in high frequency dynamics, another possibility is the use of Lagrange interpolation polynomials based on the Gauss-Lobatto quadrature points [7, 9, 10]. These functions hold the advantage that suitable mass lumping schemes are available, making it possible to fully exploit the advantages of explicit time-integration schemes. For problems with smooth solutions, the FCM achieves exponential rates of convergence, owing to the underlying high-order basis functions, without the need for a boundary-conforming mesh [3, 4, 6].

The FCM has two salient differences to the conventional finite element method. Firstly, the FCM needs a specialized approach for the numerical integration of the weak form. Due to the discontinuity introduced by the penalization of the fictitious domain, conventional Gauss-Legendre quadrature performs poorly for the FCM. Several approaches have been developed to ensure an accurate and efficient evaluation of the discontinuous domain integrals. One possibility is to use a fine grid of integration sub-cells to partition each finite cell, with the Gauss-Legendre quadrature applied for each sub-cell [3, 49]. An alternative approach applies recursive sub-division of the finite cells cut by the boundary, leading to a spacetree of integration sub-cells [6]. The spacetree approach is robust and easy to implement, but generates a high number of quadrature points. More sophisticated integration schemes have been developed to efficiently evaluate the discontinuous domain integrals for the FCM. These schemes include the blended partitioning using the smart-octree [50, 51], moment-fitting [52], adaptively weighted quadratures [53, 54], and quadratic re-parametrization for the tetrahedral finite cell method [35].

The other principal difference to conventional FEM is the enforcement of Dirichlet boundary conditions. As the boundary is no longer directly resolved by the finite elements, the Dirichlet boundary conditions cannot be enforced in the classical manner by manipulating the corresponding entries of the stiffness matrix and the load vector [55]. The are several possibilities to impose embedded Dirichlet boundary conditions in the FCM—including the penalty method, Lagrange multipliers and Nitsche's method [5, 56, 36]. This requires the additional discretization of the Dirichlet boundary Γ_D to perform numerical integration. Additionally, in case of inhomogeneous Neumann boundary conditions, a surface mesh might be needed for their numerical integration.

2.3 Weak enforcement of interface conditions

For material interface problems, the standard FCM—where a single mesh is used to discretize the entire extended domain—requires additional means to maintain its attractive approximation properties. The solution exhibits a weak discontinuity across the interface (a kink in the displacement field), and this discontinuity cannot be well represented by the high-order polynomial basis functions used in the FCM. This leads to an oscillatory solution in the cells cut by the material interface, and the convergence is reduced to a slow algebraic rate.

To overcome this challenge, we use a separate FCM mesh for each extended sub-domain:

$$\hat{\Omega}^{(i)} = \Omega^{(i)} \cup \Omega^{(i)}_{\text{fict}}, \tag{7}$$

and couple the FCM meshes to weakly impose the interface conditions, as presented in [34]. Figure 1 illustrates this approach for two sub-domains. The use of a separate set of basis functions for each sub-domain allows the discretization to approximate the strain fields with different gradients on each sub-domain. Additionally, the coupling of the FCM meshes recovers the weak discontinuity (jump in the strain field). Thereby, the FCM is able to recover its exponential convergence rates for piecewise-smooth problems [34, 36]. Moreover, this makes it possible to use discretizations with different resolutions in each sub-domain, which allows for a higher numerical efficiency. However, this multiple-mesh approach requires an additional discretization of the material interfaces Γ_{ij} for the integration of the coupling terms along the interfaces. Fortunately, the discretization of the interface for integration is considerably less demanding than FEM mesh-generation for multiple sub-domains.

Several methods can be applied for the weak coupling of the non-matching discretizations. A commonly used approach, which has also been applied in the context of the FCM, is the application of Nitsche's method [34]. It is variationally consistent, but requires the estimation of additional stabilization parameters. In the present work, we use a pure penalty approach for the coupling of FCM meshes, and for the weak enforcement of Dirichlet boundary conditions. The penalty method is inconsistent, and the accuracy depends on the choice of the penalty parameters. However, it is easy to implement—and it delivers sufficiently accurate solutions (for engineering purposes), as demonstrated by the numerical experiments presented in Section 4.

Finally, the weak form of the coupled FCM problem reads:

Find $\boldsymbol{u}^{(i)} \in \mathcal{S}(\hat{\Omega}^{(i)}), \quad i = 1, \dots, n$ such that:

$$\hat{\mathcal{B}}(\boldsymbol{u}, \boldsymbol{v}) = \hat{\mathcal{F}}(\boldsymbol{v}) \qquad \forall \boldsymbol{v}^{(i)} \in \mathcal{V}(\hat{\Omega}^{(i)}) \tag{8}$$

with:

$$\begin{aligned}
\hat{\mathcal{B}}(\boldsymbol{u}, \boldsymbol{v}) = &\sum_i \int_{\hat{\Omega}^{(i)}} \boldsymbol{\varepsilon}(\boldsymbol{v}^{(i)}) : \alpha^{(i)}(\boldsymbol{x})\ \mathbb{C}^{(i)} : \boldsymbol{\varepsilon}(\boldsymbol{u}^{(i)})\, \mathrm{d}\Omega \\
&+ \sum_{\substack{ij \\ i<j}} \beta_{ij} \int_{\Gamma_{ij}} [\![\boldsymbol{v}]\!]_{ij} \cdot [\![\boldsymbol{u}]\!]_{ij}\, \mathrm{d}\Gamma + \sum_i \beta_D^{(i)} \int_{\Gamma_D^{(i)}} \boldsymbol{v}^{(i)} \cdot \boldsymbol{u}^{(i)}\, \mathrm{d}\Gamma
\end{aligned} \tag{9a}$$

$$\begin{aligned}
\hat{\mathcal{F}}(\boldsymbol{v}) = &\sum_i \int_{\hat{\Omega}^{(i)}} \boldsymbol{v}^{(i)} \cdot \alpha^{(i)}\ \boldsymbol{b}^{(i)}(\boldsymbol{x})\, \mathrm{d}\Omega + \sum_i \int_{\Gamma_N^{(i)}} \boldsymbol{v}^{(i)} \cdot \hat{\boldsymbol{t}}^{(i)}\, \mathrm{d}\Gamma \\
&+ \sum_i \beta_D^{(i)} \int_{\Gamma_D^{(i)}} \boldsymbol{v}^{(i)} \cdot \hat{\boldsymbol{u}}^{(i)}\, \mathrm{d}\Gamma
\end{aligned} \tag{9b}$$

where $\boldsymbol{v}^{(i)}$ are the test functions in the space of admissible functions $\mathcal{V}(\hat{\Omega}^{(i)})$, and

$$\alpha^{(i)}(\boldsymbol{x}) = \begin{cases} 1 & \forall \boldsymbol{x} \in \Omega^{(i)} \\ \epsilon^{(i)} \ll 1 & \forall \boldsymbol{x} \in \Omega^{(i)}_{\text{fict}}. \end{cases} \tag{10}$$

The two penalty parameters, β_{ij} and β_D control the weak enforcement of the kinematic compatibility condition across Γ_{ij} and of the Dirichlet boundary condition, respectively. Here, $[\![\boldsymbol{u}]\!]_{ij}$ denotes the jump operator

$$[\![\boldsymbol{u}]\!]_{ij} = \boldsymbol{u}^{(i)} - \boldsymbol{u}^{(j)}. \tag{11}$$

The choice of the penalty parameters influences the accuracy of the numerical solution. Values that are too small lead to a large error in the imposed conditions (observed in the levelling-off of the convergence curves), whereas values that are too large yield an ill-conditioned system of equations [57–59]. In the numerical examples that follow in this paper, the penalty parameters are chosen heuristically, starting with an initial value of $E \cdot 10^3$, where E is the Young's modulus of the stiffest material. In case the resulting error in the imposed condition is too large, the penalty value is increased stepwise by one order of magnitude.

3 Multi-level *hp*-refinement

For problems with local solution characteristics such as stress concentration and singularities, *hp*-refinement is a powerful method of discretization, as it is more efficient than uniform *h*- or *p*-refinement [40–42]. Smaller elements with low polynomial degree close to the singularities are able to capture the local solution and confine the discretization error, whereas large elements with high polynomial degree are best suited to describe the smooth solution further away from singularities.

However, the Cartesian grid nature of the FCM discretization does not easily incorporate local mesh refinement without having to introduce hanging nodes. In this work, we apply the multi-level *hp*-refinement scheme to locally refine FCM grids, without the difficulties of constraining hanging nodes. Thereby, the method is able to achieve higher convergence rates, while retaining the advantages of using Cartesian finite cell grids. This section provides a brief review of the multi-level *hp*-FEM and elaborates on the refinement strategy used for material-interface problems.

3.1 Basic refinement concept

The principal idea of the multi-level *hp*-FEM is to enhance the local approximation of the solution through a local superposition of coarse base elements in the refinement zones with finer overlay elements. The final approximation $\boldsymbol{u}$ is the sum of the base mesh solution $\boldsymbol{u}_b$ and the fine-scale overlay solution $\boldsymbol{u}_o$. This *refine-by-superposition* approach is related to the pioneering work of Mote [60]. This concept was applied in the context of *hp*-domain decomposition, to overlay coarse *p*-FEM meshes with fine linear *h*-elements [61], and in the hierarchic *hp-d*-scheme, which adds several layers of *linear* overlay meshes [62].

The multi-level *hp*-scheme generalizes the hierarchic overlay idea of the *hp-d* scheme by introducing high-order overlay meshes. Here, the support of the high-order basis functions is chosen on the finer overlay meshes—instead of the base discretization, as illustrated in Figure 2. This corresponds to an *h*-refinement of the high-order basis functions, as the size of their support is reduced. This limited support—combined with an adequate geometric grading—has the desirable effect of confining the (global) pollution error in the smallest elements around the singularities. This is not the case for the linear overlay approach presented in [61] and the hierarchic *hp-d* linear overlays [62], as the error propagates in the large base elements of the base mesh [45]. Consequently, the multi-level *hp*-FEM leads to superior convergence rates for problems with singularities and sharp features. Numerical experiments demonstrate that the multi-level *hp*-FEM converges exponentially, for problems including vertex and edge singularities in two and three dimensions [45, 47, 46].

The multi-level *hp*-FEM avoids the implementational burden of constraining hanging nodes. A

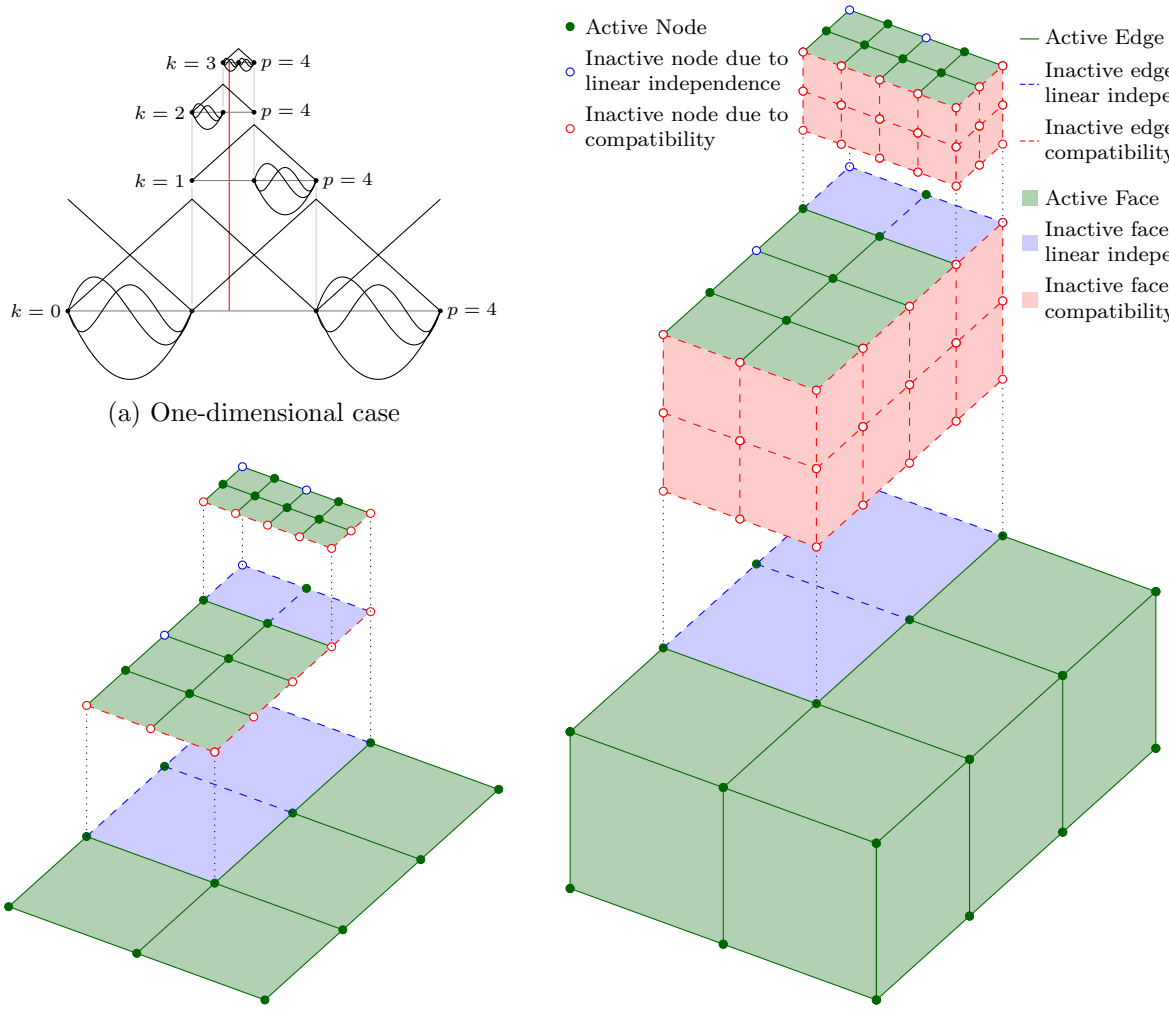

(a) One-dimensional case

(b) Two-dimensional case

(c) Three-dimensional case

Figure 2: Conceptual idea of the multi-level *hp*-method [47].

simple implementation is possible, wherein two requirements are satisfied to ensure convergence: *compatibility* of the discretization, and *linear independence* of the basis functions. The overlay meshes, generated by recursive sub-divisions of the base mesh, are geometrically incompatible. To guarantee compatibility of the discretization, homogeneous Dirichlet boundary conditions are prescribed on the boundary of each layer of overlay meshes. This maintains a C^0-continuous approximation by construction. To ensure linear independence of the basis functions, the high-order shape functions on overlay elements are excluded from their respective parent elements, as demonstrated in Figure 2. The fulfillment of these requirements leads to a simple yet powerful discretization technique, that exploits the benefits of the *hp*-FEM without the burden of constraining hanging nodes. Additionally, using a suitable data structure in an object-oriented framework allows for a dynamic update of the overlay meshes without introducing a noticeable computational cost [47]. For an in-depth discussion of the multi-level *hp*-FEM for two- and three-dimensional problems, with elaboration of their implementational aspects, the interested reader is referred to [45, 47, 46]

3.2 Refinement strategy

In general, stress concentration and sharp solution features can be caused either by sudden changes in the boundary conditions, the applied body load, or sharp geometric features—such as re-entrant corners. In this work, we mainly consider stress concentrations induced by sharp geometric features of the material interfaces. To resolve such features, an *a priori* refinement scheme is applied. Here,

the design of the discretization is guided by existing knowledge of the solution features, based on engineering experience.

3.2.1 *h*-refinement

Typically, *hp*-mesh design includes a geometric progression of elements towards the singularities, with small linear elements close to the singularities, and increasing element size as well as increasing polynomial orders away from the singularities. In the framework of the multi-level *hp*-FEM, the high-order overlay meshes are generated by recursive bisection of the base mesh. To geometrically grade the mesh towards a point $\boldsymbol{P_x} = (p_x, p_y, p_z)$, for each leaf element K, the point is mapped into the index space of the element $\Omega_\square = [-1, 1]^d$ as follows:

$$\boldsymbol{P}_{\boldsymbol{\xi}_K} = \boldsymbol{Q}_K^{-1}(\boldsymbol{P_x}), \tag{12}$$

where $\boldsymbol{P}_{\boldsymbol{\xi}_K} = (p_\xi, p_\eta, p_\zeta)$ are the local coordinates of $\boldsymbol{P_x}$ in the index space of the leaf-element K, and $\boldsymbol{Q}_K(\boldsymbol{x})$ denotes the geometric mapping between the global and local spaces. The element is sub-divided if the mapped point lies within an extended index space $\hat{\Omega}_\square$, which is larger than $\Omega_\square$

$$\boldsymbol{P}_{\boldsymbol{\xi}_K} \in \hat{\Omega}_\square. \tag{13}$$

Typical choices for $\hat{\Omega}_\square$ are between $[-1.25, 1.25]^d$ and $[-2, 2]^d$. Thereby, not only the elements containing $\boldsymbol{P_x}$ are refined, but possibly also the neighboring elements. By carrying out several recursive refinement steps, this approach produces a geometrically graded mesh—with decreasing element size towards the point of interest, as demonstrated in Figure 3. For Cartesian grid meshes, which are used in the FCM, this inverse mapping carries a trivial computational cost.

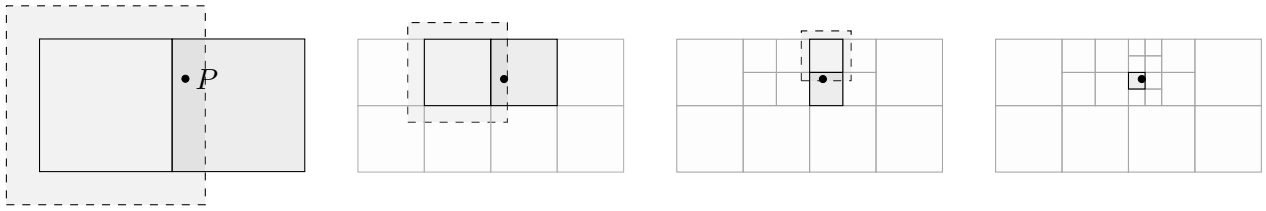

(a) $k = 0$ (b) $k = 1$ (c) $k = 2$ (d) $k = 3$

Figure 3: Mesh grading towards $\boldsymbol{P}$ using an extended index space $\hat{\Omega}_\square = [-1.25, 1.25]^2$. An element is refined if $\boldsymbol{P}$ lies within its extended space (drawn with dashed lines).

3.2.2 *p*-distribution

The distribution of the polynomial degree p across the layers determines the convergence properties of the *hp*-discretization. The optimum p-distribution is not known *a priori*.

In one case, the high-order basis functions can be kept on the base mesh, and only linear basis functions are used on the overlay meshes (Figure 4a). This case corresponds to the hierarchical *hp-d* scheme [62], as a special case of the multi-level *hp*-scheme. The linear overlay meshes introduce a moderate number of additional unknowns. However, the large support of the high-order basis functions causes the global error to propagate through the domain, yielding a minimal improvement in the convergence rates [45].

On the other end of the spectrum, the high-order basis functions are all shifted to the highest (finest) overlay mesh (Figure 4b). Obviously, this uniform multi-level *hp*-approach introduces significantly more unknowns. However, with the limited support of the high-order functions, the discretization error is confined to the smallest elements, which yields superior convergence properties.

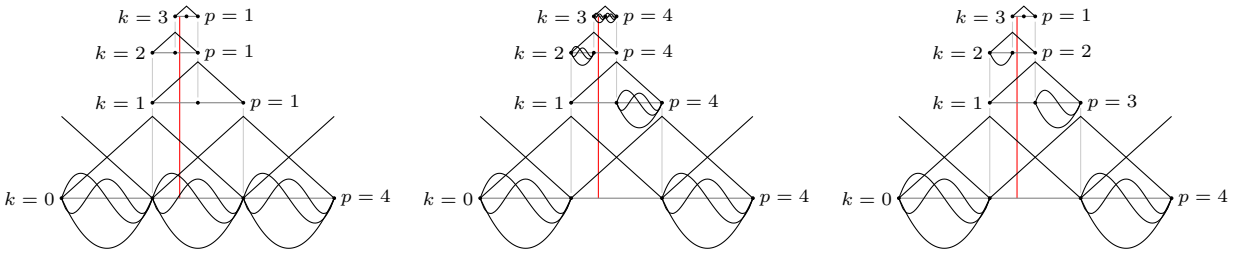

(a) Linear overlay, *hp-d*-approach (b) Uniform multi-level *hp*-approach (c) Non-uniform multi-level *hp*-approach

Figure 4: Comparison of different approaches for hierarchical, high-order refinement [47].

In this work, we adopt the uniform multi-level *hp* refinement approach, which is equivalent to uniform *p*-elevation on a geometrically graded mesh. This scheme has been shown to be robust and efficient, although it might not be optimal regarding the approximation error for a number of degrees of freedom. For singular problems, the uniform multi-level *hp*-refinement shows exponential convergence in the pre-asymptotic range, which can be extended by increasing the refinement depth [45, 47]. For boundary conforming problems, an *a posteriori* error estimator for multi-level *hp*-FEM coupled with a smoothness indicator was developed [63], where automatically generated multi-level *hp*-FEM discretizations with non-uniform *p*-distributions further improved the efficiency.

4 Numerical experiments

In this section, we examine a series of benchmarks involving material interfaces where the geometry induces high solution gradients and singularities. This includes the two-dimensional benchmark of a plate with an elliptical inclusion, where we perform a systematic study of the influence of the numerical integration accuracy on the quality of the approximation. The second benchmark considers a circular plate with a sharp inclusion, where the kink in the material interface introduces a vertex singularity. In three dimensions, we examine the problem of an embedded ellipsoidal inclusion, as well as a singular benchmark of a cylinder with an embedded cubical inclusion, which causes a vertex-edge singularity. In all examples, we examine the *p*-convergence and the numerical approximation of the first derivatives of the solution.

4.1 Plate with an elliptical inclusion

4.1.1 Problem setup

The first material interface benchmark we study is a plane stress problem of a plate $\Omega^{(1)}$ with an embedded elliptical inclusion $\Omega^{(2)}$, which has a large aspect ratio, $r_y/r_x = 0.15$, and a moderate stiffness ratio to the plate, $E^{(2)}/E^{(1)} = 10$ (see Figure 5a). The displacement field within the plate is distorted by the curved inclusion, leading to high stress concentration in the plate. The stress concentration is amplified by the high curvature of the interface at the major axis of the ellipse (points A and C) as shown in Figures 5c and 5b. The displacement field within the elliptic inclusion remains smooth. The exact solution for the displacements is piecewise-analytic, i.e. it can be *exactly* represented by a Taylor expansion in each sub-domain. An overkill solution was obtained for reference using a conforming *p*-FEM mesh which has 26 elements with a uniform polynomial degree $p = 30$ and 47,102 degrees of freedom. In the conforming mesh, the elements at the interface are blended on the ellipse's exact geometry using a quasi-regional mapping [64]. For the setup shown here, the strain energy obtained using the overkill discretization is

$$\mathcal{U}_{\text{ex}} = 9.10131116644 \times 10^{-2}. \tag{14}$$

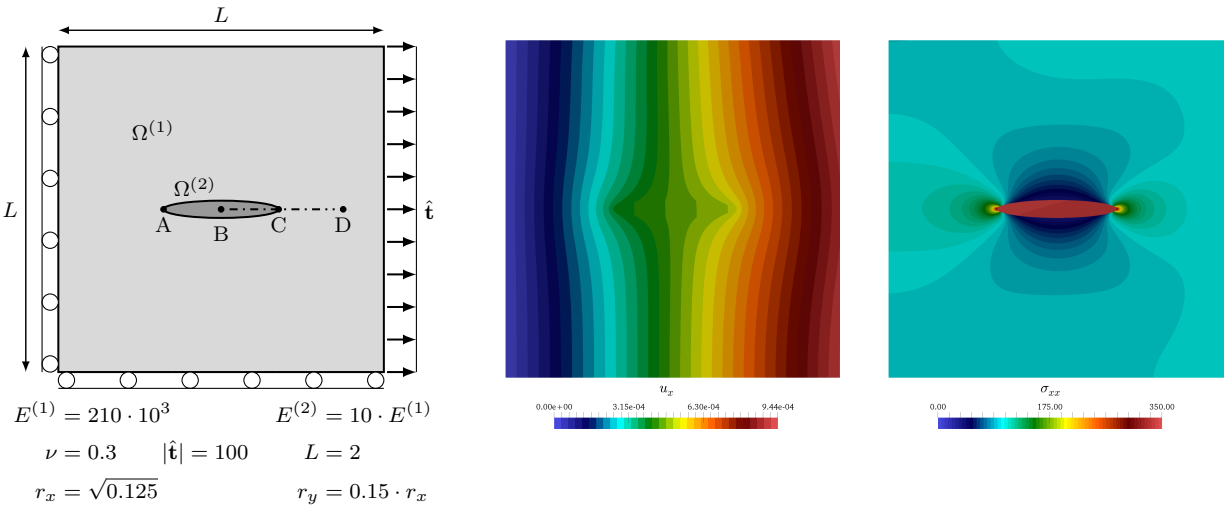

(a) Benchmark setup (b) Reference u_x displacement (c) Reference σ_{xx} stress

Figure 5: Plate with elliptical inclusion

4.1.2 Discretization

The problem was solved numerically based on the FCM and a weak enforcement of the interface conditions. Two separate FCM meshes were used to discretize the plate and the inclusion. A base discretization matching the outer boundary of the plate with 4×4 finite cells was used. The same parametrization was used for the inclusion's mesh, excluding cells that do not intersect the ellipse's domain, leaving 2×2 finite cells—as shown in Figure 6. The ellipse's dimensions and the mesh parameters are setup such that the mesh does not exactly match points A or C for any number of recursive bisection steps. To adequately describe the local solution in the plate, the mesh was locally refined using the uniform multi-level *hp*-refinement scheme. The mesh was refined *a priori*, with a geometrically driven grading towards points A and C, as described in Section 3.2. The FCM mesh of the inclusion was left without local refinement, as the solution within the inclusion $\Omega^{(2)}$ is smooth. The use of different mesh resolutions for the two sub-domains is possible thanks to the weak coupling approach, and leads to an improved numerical efficiency, as the refinement stays local to the region of interest. To avoid ill-conditioning of the stiffness matrices, the fictitious domain in both meshes was assigned a low stiffness of $E_{\text{fict}} = E^{(1)} \cdot 10^{-11}, \nu_{\text{fict}} = 0.3$. Since the exterior boundary of the plate's mesh conforms to the true geometry, the Dirichlet and Neumann boundary conditions were enforced in a conventional manner.

For the numerical integration of the finite cells, two possibilities were considered. The first approach employs a recursive sub-division algorithm for integration sub-cells that are intersected by the geometric boundary, resulting in a spacetree—as shown in [6], for example. The second approach considered is the blended partitioning following [50], where the cut cells are automatically partitioned into boundary-conforming quadrilaterals and triangles, which are blended on the exact geometry. It needs to be pointed out that this approach produces geometry-conforming sub-cells to be used for numerical integration. Yet, it cannot be used in general to generate an analysis-suitable FEM mesh, as it does not need to satisfy requirements on shape and regularity at cell edges and faces, e.g. at the transition to the neighboring cells. This approach shares some similarities with the method presented by Fries *et al.* [65–67]. The integration meshes for both considered approaches are displayed in Figure 6.

The penalty method was used to weakly couple the two meshes at the material interface as described in Section 2.3. The penalty parameter $\beta_{12} = 10^7$ was chosen empirically. For the numerical integration of the interface coupling terms, a linear discretization of the embedded interface was generated using a marching squares algorithm[68]. The marching squares implementation uses the implicit geometric description of the ellipse and a grid of $n_g \times n_g$ points per sub-cell (leaf element) of

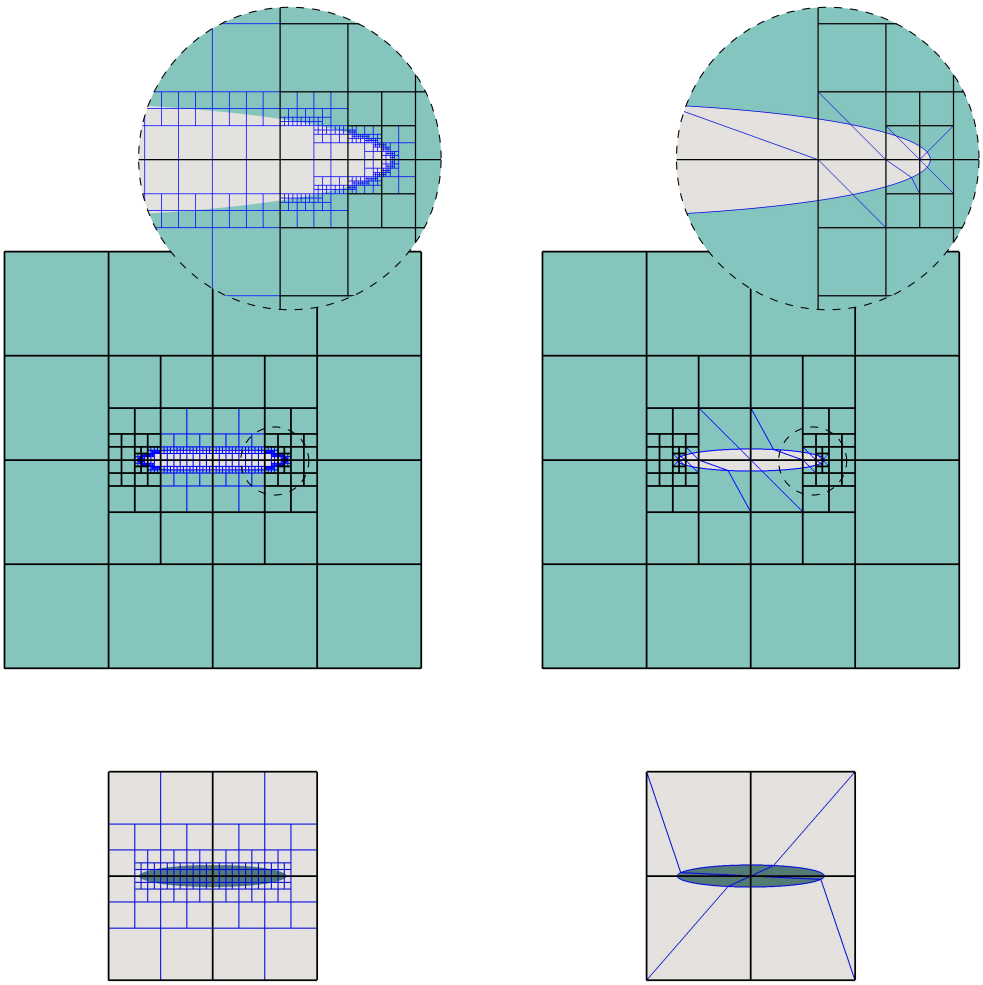

(a) Spacetree partitioning, $m = 4$, following [62]

(b) Blended partitioning, following [50]

Figure 6: Domain discretization for the plate and the inclusion, refinement depth $k = 4$. The finite cells' edges are drawn with thick black lines, whereas the integration meshes are drawn with thin blue lines.

the plate's mesh, yielding a total of n_{seg} linear segments, as can be seen in Figure 7a. An advantage of this approach is that the resulting segments do not cross the finite-cell boundaries, resulting in a more accurate evaluation of the interface integral, which the terms of which are only C^0-continuous. Additionally, the segments have the same size relative to the leaf element in which they lie. Consequently, the non-uniform resolution of the FCM mesh carries over to the discretization of the interface.

An extension to this approach generates high-order segments using Lagrange interpolation polynomials with degree p_{seg}, by adding $(p_{\text{seg}} - 1)$ interpolation nodes within the linear segments. The implicit geometry description was used in conjunction with a line-search method to correctly place the additional interpolation points on the interface. This approach finally yields a high-order boundary discretization, as shown in Figure 7b (note that the figure shows only the extreme nodes). For a more detailed description of this high-order interface recovery algorithm, the reader is referred to [69]. This convenient approach for embedded boundary parametrization highlights another advantage of using Cartesian grids for the finite cell meshes and the same mesh parameters for both domains - in contrast to using unstructured FCM meshes, where the generation of cell-conforming boundary segments would be more challenging.

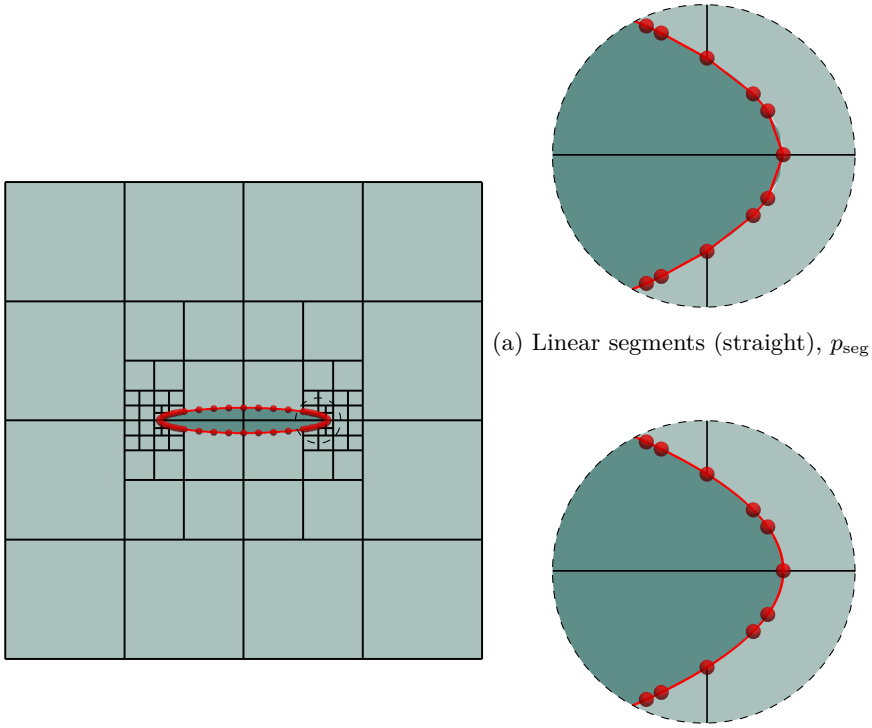

(a) Linear segments (straight), $p_{\text{seg}} = 1$

(b) Quadratic segments (curved), $p_{\text{seg}} = 2$

Figure 7: Discretization of the elliptical interface, $n_g = 4$, $n_{\text{seg}} = 72$

4.1.3 Numerical solution

Figure 8 shows the numerical approximation of the axial stresses σ_{xx} obtained using the FCM with weak coupling and multi-level *hp*-refinement. The plots shown here are the superposition of the solutions from the physical part of each domain, whereas the solution in the fictitious domain is omitted. Without applying any local refinement (Figure 8a), the FCM with weak coupling is able to reproduce the overall solution characteristics. However, the approximation for the stresses exhibits severe oscillations. In particular, the four finite cells in the plate's mesh that are intersected by the interface are affected by the high curvature of the concave geometry at points A and C, leading to stress concentration. This local solution characteristic cannot be represented well by the smooth basis functions of the coarse mesh, which leads to an oscillatory approximation of the stresses. Figure 8b depicts the numerical approximation obtained with the multi-level *hp*-scheme, using two levels of overlay meshes with uniform *p*-distribution. Applying local refinement to the FCM mesh of the plate improves the numerical approximation significantly. The solution benefits from the local support of the high-order basis functions with C^0-continuity, and the oscillations in the stresses are confined to the cells on the finest overlay mesh. Increasing the refinement depth with $k = 4$, the quality of the approximation improves further.

Figure 9 examines the stresses along the cutting line B-D (see Figure 5a) in closer detail. Since the cutting line is orthogonal to the ellipse's boundary, the exact solution for σ_{xx} is continuous along the cutting line. Hence, jumps in the numerical approximation of the stresses, either at the interface or at element boundaries, correspond to the approximation error directly. The numerical approximation obtained without local refinement exhibits severe oscillations and a large jump at the interface. Due to the coupling condition, the oscillations extend into both domains. Due to the oscillations, the maximum stresses are overestimated and the local stress distribution is misrepresented. Applying multi-level *hp*-refinement to the plate's FCM mesh enhances the quality of the approximation appreciably, as the oscillations are mainly restricted to the finite cells on the finest overlay mesh. The zoomed plot at the interface reveals the fine-scale features of the solution in the plate. The local solution characteristics are not fully resolved by the refined mesh with $k = 4$, which suggests that this problem would benefit from further local refinement.

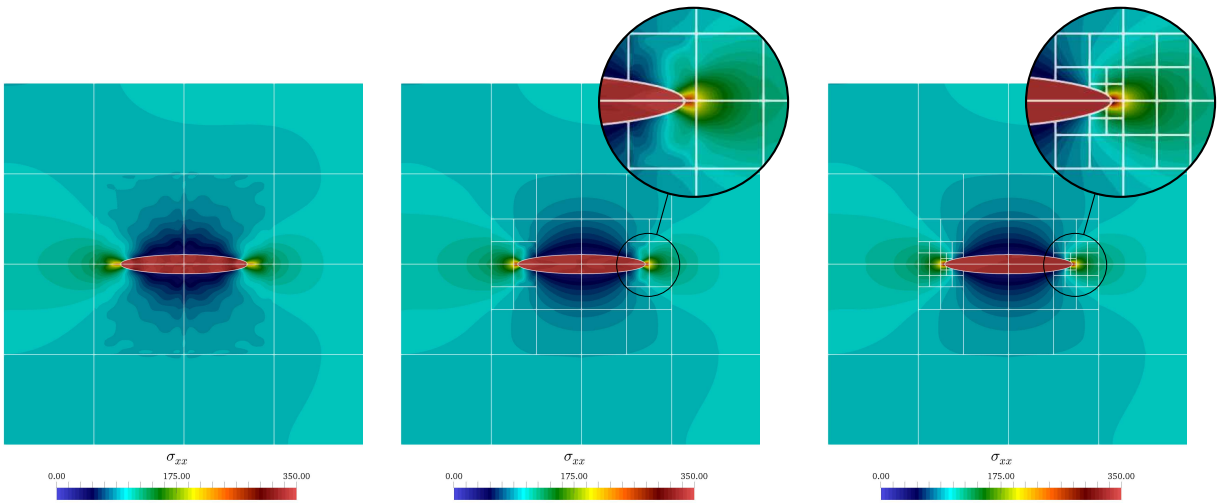

(a) $p = 8$, no refinement (b) $p = 4, k = 2$ (c) $p = 4, k = 4$

Figure 8: Plate with elliptical inclusion: numerical approximation of the axial stresses σ_{xx}

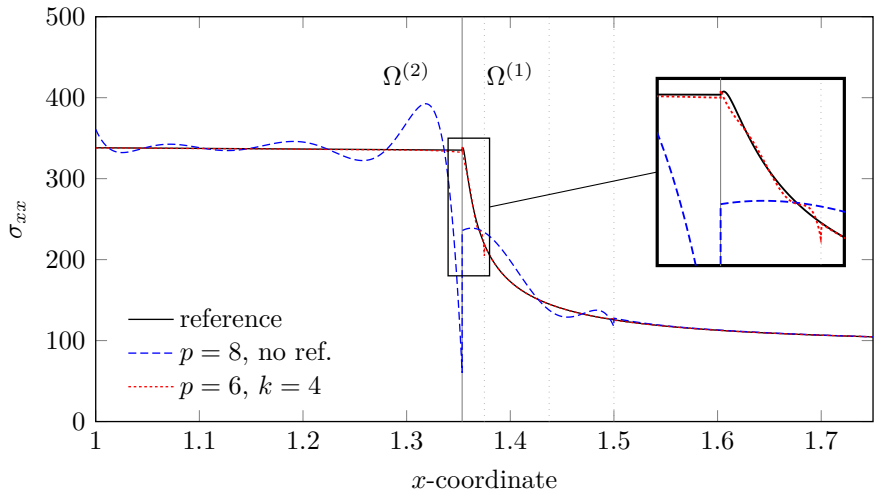

Figure 9: Plate with elliptical inclusion: numerical approximation of the axial stresses σ_{xx} along cutting line B-D

4.1.4 Influence of domain integral accuracy

To assess the accuracy of the numerical integration for the domain integral, Figure 10 compares the error in σ_{xx} along cutting line B-D, based on multi-level hp-refinement of the plate's mesh with $k = 4$, polynomial degree $p = 6$, and different domain partitioning schemes. The error here is calculated with respect to the reference solution generated using the overkill p-FEM discretization. The blended partitioning approach gives a highly accurate evaluation of the domain integrals, as it uses the exact geometric description of the interface. Using the spacetree partitioning scheme with a depth $m = 2$ leads to high approximation errors. The error is highest in the finite cell from inclusion's mesh, which is attributed to the coarser spatial resolution of the integration, as demonstrated by Figure 6a. The high error propagates through the mesh, carrying over to finite cells which are not cut by the material interface, and to the other mesh through the coupling condition. Increasing the depth m of the spacetree reduces the approximation error appreciably. Using depth $m = 7$, the achieved accuracy is comparable to that of the blended partitioning approach.

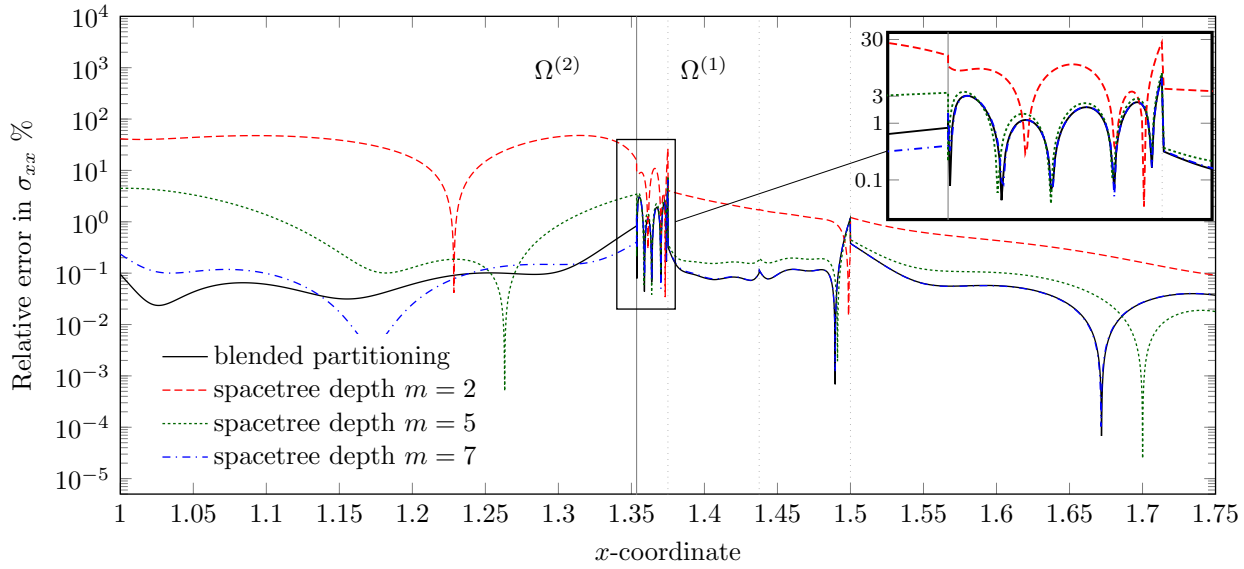

Figure 10: Error in σ_{xx} along cutting line B-D with different domain integration schemes, $p = 6$, $k = 4$

4.1.5 Influence of interface integral accuracy

To evaluate the effect of the integration accuracy of the interface integral, Figure 11 plots the local distribution of the approximation error for the von-Mises stresses σ_{eq} along the interface in the plate. Using a grid dimension of $n_g = 4$ for the marching squares implementation, $n_{\text{seg}} = 72$ linear segments are generated (see Figure 7a). However, this linear discretization of the curved interface introduces artificial stress singularities at the interpolation nodes between the segments, as the segments touch at the extreme nodes but have discontinuous normals, which introduces artificial kinks at the interpolation nodes. At ellipse parameter $t = 0$ (point C), the artificial singularity takes over the approximation error, causing the entire solution to diverge. To reduce the error at the interface, h-refinement of the interface's discretization is carried out by refining the grid for the marching squares. Figure 11a shows the resulting error distribution with $n_g = \{4, 8, 16, 32, 64\}$, which yields $n_{\text{seg}} = \{72, 136, 280, 564, 1120\}$. The h-refinement of the segments delivers a uniform decrease in the approximation error, with a maximum error in σ_{eq} of approximately 3% using 1120 segments.

Alternatively, p-refinement of the interface segments using Lagrange polynomials, as described earlier, leads to a faster convergence as demonstrated by Figure 11b. Using 72 quadratic segments already decreases the error significantly, as the high-order discretization approximates the tangents at the interpolation nodes with better accuracy, as shown in Figure 7b. Increasing p_{seg} further from 3 to 4 in this example, gives a marginal improvement in accuracy, indicating that the numerical integral at the interface converges. The remaining approximation error is attributed to the FCM discretization with $p = 8$, and $k = 6$. The two refinement schemes for the interface discretization converge to the same local error distribution.

4.1.6 Convergence study

A uniform p-refinement study was carried out to assess the overall convergence behavior of multi-level hp-refinement in conjunction with weak coupling. Figure 12a shows the relative error in the energy norm for different refinement depths

$$||e||_E = \sqrt{\frac{|\mathcal{U}_{\text{num}} - \mathcal{U}_{\text{ex}}|}{\mathcal{U}_{\text{ex}}}} \cdot 100\%, \tag{15}$$

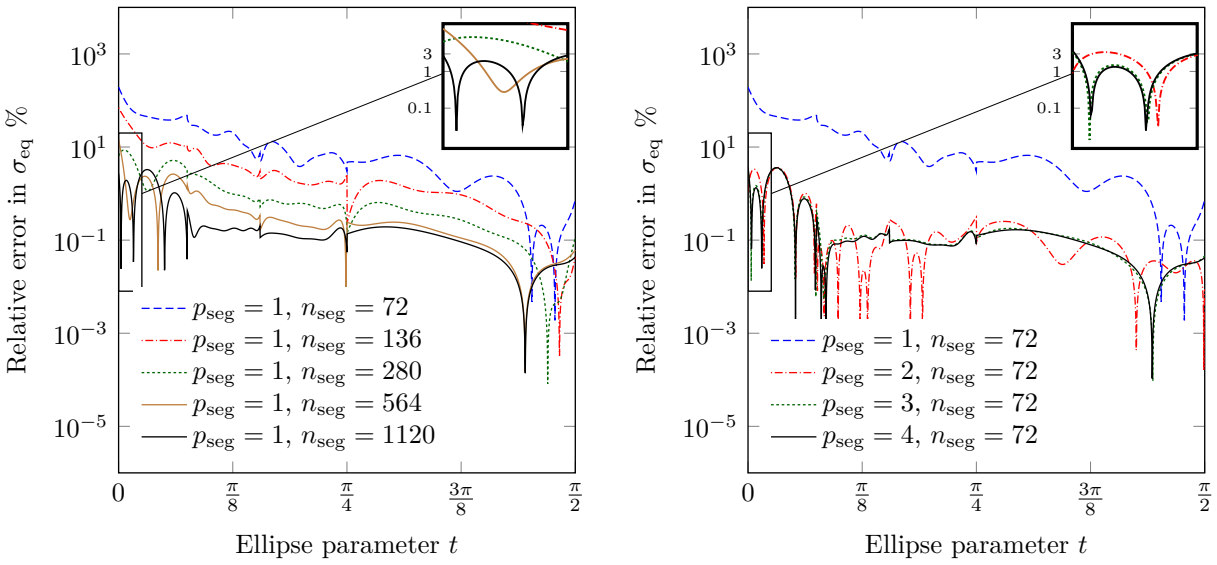

(a) h-refinement

(b) p-refinement

Figure 11: Error in σ_{eq} along the interface using different interface discretizations, $p = 8$, $k = 4$

where $\mathcal{U}_{\text{num}}$ denotes the numerically approximated strain energy, and $\mathcal{U}_{\text{ex}}$ denotes the reference strain energy. For this study, the blended approach is used for the domain integration, in conjunction with a high-order parametrization of the interface. As the exact solution is piecewise-analytic, FCM with a weak enforcement of the interface conditions converges exponentially under uniform p-elevation, even without applying any local refinement. For a more detailed analysis of the convergence rates, Figure 12b shows the convergence plot in a log-$\sqrt{\ }$ scaling. Here, the linear plots characterize exponential convergence in the form [70]:

$$||e||_E \leq C \exp\left(\gamma N^{\theta}\right) \quad \text{with } \theta = \frac{1}{2}, \tag{16}$$

where C is a positive constant, and γ is a negative constant describing the convergence rate. Increasing the refinement depth with the multi-level hp-scheme for this example leads to a steeper convergence rate γ, allowing for a higher gain in accuracy. To achieve an engineering accuracy of 1% error, the discretization with $k = 4$ needs approximately five times less the degrees of freedom than the non-refined case.

4.2 Bi-material inclusion corner

In the next study, we consider a singular material-interface benchmark problem adapted from [71]. The two-dimensional temperature problem is governed by the Poisson equation:

$$\kappa^{(i)} \nabla^2 \phi^{(i)} = -1 \qquad \forall \boldsymbol{x} \in \Omega^{(i)} \tag{17}$$

$$\phi = 0 \qquad \forall \boldsymbol{x} \in \Gamma_D, \tag{18}$$

where $\phi^{(i)}$ denotes the temperature, $\kappa^{(i)}$ the thermal diffusivity, $\Omega^{(i)}$ and Γ_D are defined as shown in Figure 13a. Here, the material interface Γ_{12} has a sharp corner, inducing a vertex singularity. Moreover, the intersection of the material interface with the Dirichlet boundary Γ_D for the applied boundary condition introduces two additional weak singularities, where the solution exhibits reduced continuity.

The exact solution to this problem is given in radial coordinates (r, θ) by [72] :

$$\phi(r, \theta) = A_1 r^{\lambda_1} h_1(\theta) + A_2 r^{\lambda_2} h_2(\theta) + \mathcal{O}(r^2), \tag{19}$$

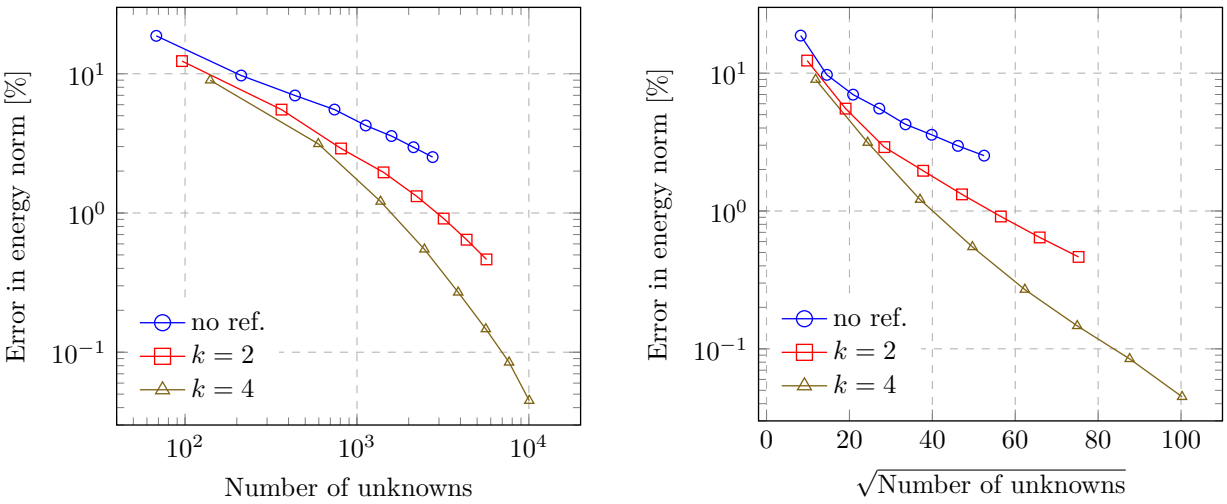

(a) Convergence in log-log scale

(b) Convergence in semi-log scale

Figure 12: Plate with elliptical inclusion: p-convergence for multi-level hp-refinement

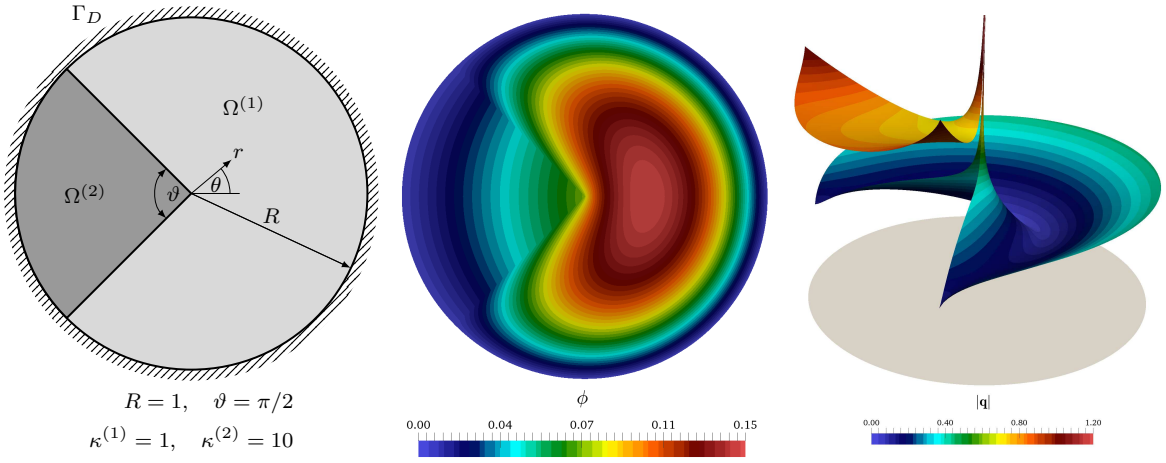

(a) Benchmark setup

(b) Reference temperature ϕ

(c) Reference heat flux $|\boldsymbol{q}|$

Figure 13: Bi-material inclusion corner

where A_1 and A_2 are scalar constants, $h_1(\theta)$ and $h_2(\theta)$ are smooth sinusoidal functions and

$$\lambda_1 = 0.731691779, \quad \lambda_2 = 1.268308221. \tag{20}$$

For reference, an overkill solution was generated using a conforming p-FEM with exact blending, where the mesh was geometrically graded towards the vertex singularity. The reference mesh has 21 elements with polynomial degree $p = 30$ and 18,991 degrees of freedom. The strain energy for the setup shown here is

$$\mathcal{U}_{\text{ex}} = 1.0168443145 \times 10^{-1}. \tag{21}$$

Using the FCM with weak coupling, separate meshes were used for the disc and the inclusion, as shown in Figure 14. Unlike the previous example, where the stress concentration in the plate was caused by the concave geometry, the singularities affect the solution in both sub-domains. Consequently, *both meshes* need local refinement to resolve the singularities. The meshes were graded towards the vertex singularity, and the two weak singularities on the boundary. The same refinement depth was applied to all three refinement points. Here, the discretizations match the

singularities with nodes of the FCM mesh.

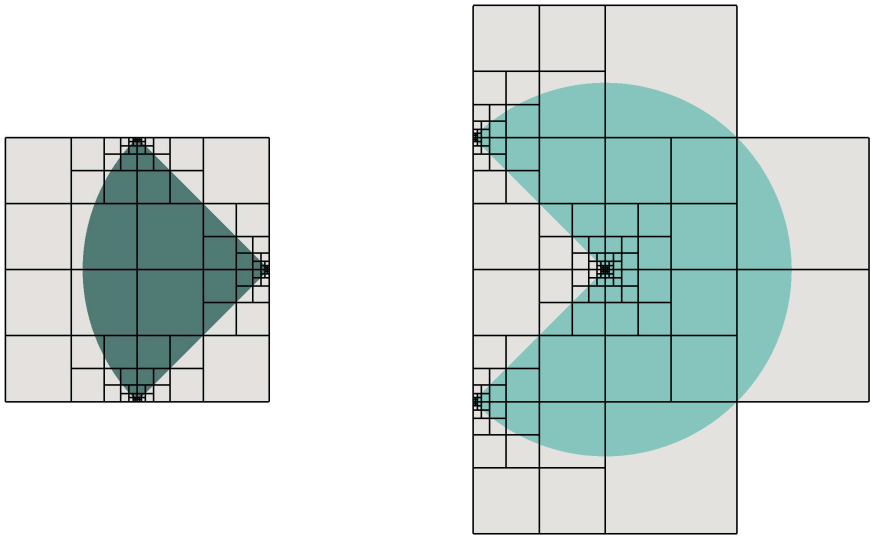

(a) Inclusion (b) Disc

Figure 14: Bi-material inclusion corner: discretization with refinement depth $k = 6$

The penalty method was used to couple the FCM meshes along the straight material interface. Additionally, the embedded Dirichlet boundary condition was weakly enforced on $\Gamma_D^{(i)}$ in each mesh using the penalty method with $\beta_D = 10^7$. The weak enforcement of the boundary conditions does not influence the convergence behaviour, provided that the penalty parameter is chosen large enough [58]. The fictitious domain was penalized with $\kappa_{\text{fict}} = \kappa^{(1)} \cdot 10^{-9}$. Note that this value is higher than in the previous example, and was chosen to have a lower condition number to allow the use of high polynomial degrees, up to $p = 10$, in the convergence studies. The blended partitioning approach was used to numerically evaluate the domain integral. The coupling integral was evaluated on linear segments, whereas the weak Dirichlet boundary conditions were integrated over curved segments, automatically generated with the same approach described for the previous example.

The magnitude of numerical solution for the heat flux, $\boldsymbol{q}^{(i)} = \kappa^{(i)} \nabla \phi^{(i)}$, is depicted in Figure 15. For the first case, without any local refinement of the meshes, the numerical approximations of the heat fluxes exhibit oscillations within the domain and jumps at element boundaries. The oscillatory behavior is attributed to the inability of the smooth polynomials to represent the high gradients in the vicinity of the singularities. For the second discretization, two levels of uniform high-order overlay meshes are used to grade the meshes towards the three singularities.The C^0-continuity of the high-order basis functions limits the pollution error from the singularities at the finest refinement level, as revealed in the zoomed plot. Using four refinement levels improves the solution further, as the geometric progression of element size arrests the propagation of the error, while having larger elements to describe the smooth solution away from the singularities.

Figure 16a shows the results of a p-elevation study carried out to assess the convergence properties. Without applying local refinement, the convergence under p-elevation is identified as being algebraic in the form

$$||e||_E \leq aN^{\eta}, \tag{22}$$

where a is a positive constant, and η is a negative constant indicating the convergence rate. The estimate for $\eta \approx -\lambda_1$ matches the theoretical expectation for a discretization where the vertex singularities are matched by nodes of the mesh [73]. With multi-level hp-refinement the convergence behavior shows a pre-asymptotic range with a higher convergence rate, and an asymptotic range, where the behavior returns to being algebraic with the same rate η. Increasing the refinement depth k extends the pre-asymptotic range, allowing for a significant decrease in the total approximation

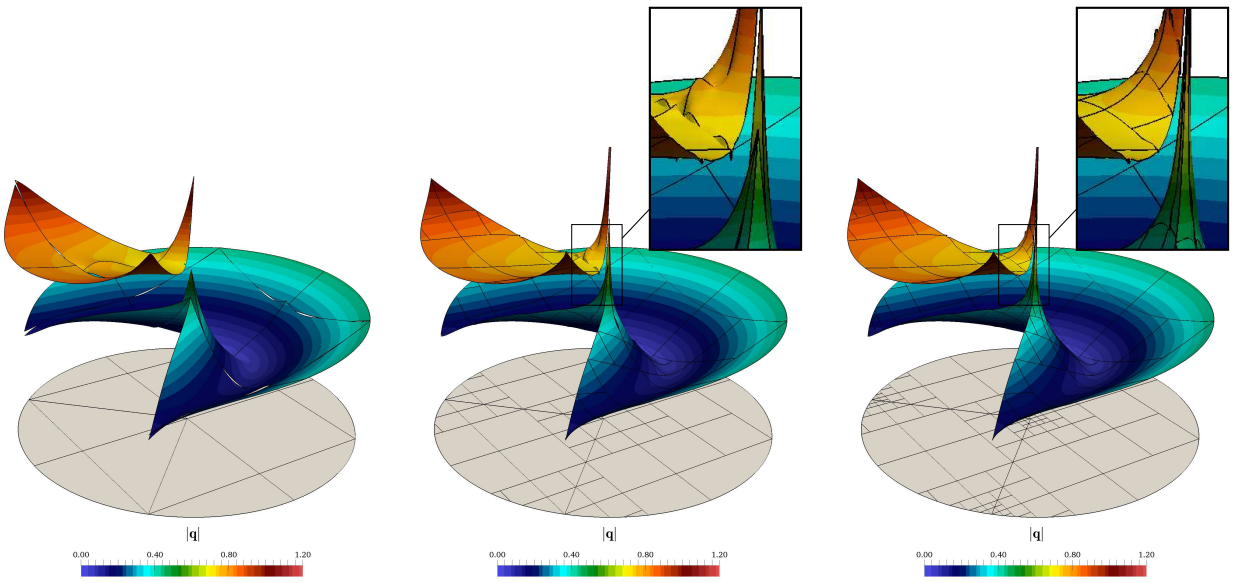

(a) $p = 4$, no refinement (b) $p = 4, k = 2$ (c) $p = 4, k = 4$

Figure 15: Bi-material inclusion corner: numerical approximation of the flux $|\boldsymbol{q}|$

error. Changing the scaling to a log-$\sqrt[3]{\cdot}$ in Figure 16b, the extended pre-asymptotic range appears linear, indicating exponential convergence in the pre-asymptotic range.

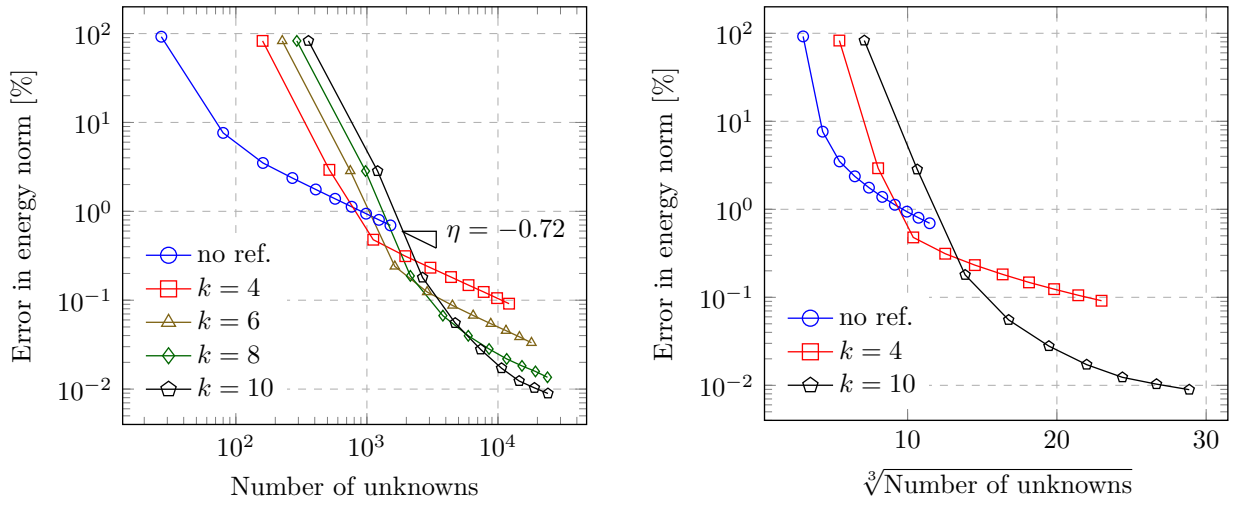

(a) Convergence in log-log scale (b) Convergence in semi-log scale

Figure 16: Bi-material inclusion corner: p-convergence for multi-level hp-refinement

Note the apparent leveling-off of the error in the energy norm around $10^{-2}\%$, which is attributed to the use of $\kappa_{\text{fict}} \neq 0$, adding a modeling error to the numerical approximation. The introduced error in the energy norm is in the order of $\sqrt{\kappa_{\text{fict}}}$ [74]. Using a smaller value for κ_{fict} lowers the leveling-off threshold, at the cost of the conditioning of the system.

4.3 Cube with ellipsoidal inclusion

Next, we consider a three-dimensional example, which is a generalization of the benchmark presented in Section 4.1. A linear elastic analysis is considered for a cube with an embedded ellipsoidal

(prolate spheroid) inclusion - depicted in Figure 17. The spheroidal inclusion induces stress concentration in the cube around the two points on its major axis. The solution within the spheroidal inclusion remains smooth.

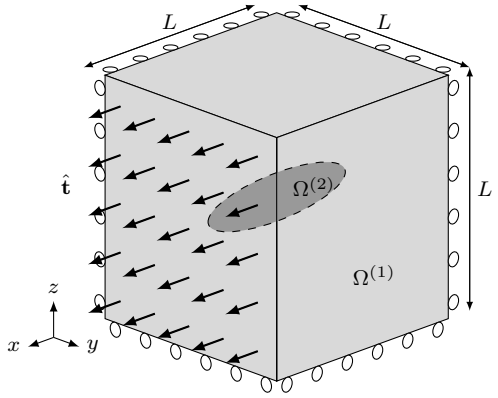

$$E^{(1)} = 210 \cdot 10^3, \quad E^{(2)} = E^{(1)} \cdot 10,$$
$$\nu = 0.3, \quad |\hat{\mathbf{t}}| = 100, \quad L = 1,$$
$$r_x = \sqrt{0.125}, \quad r_y = r_z = 0.25 \cdot r_x$$

Figure 17: Cube with ellipsoidal inclusion

Two FCM meshes were used to discretize the sub-domains, using a base discretization of $4 \times 4 \times 4$ finite cells which match the outer boundaries of the cube. To resolve the stress concentration, the cube's mesh was refined towards the two points on the major axis of the spheroid, as shown in Figure 18. For the inclusion's mesh, $2 \times 2 \times 2$ elements remain after excluding all cells that are completely in the fictitious domain. Finite cells that are cut by the interface were further partitioned for integration using the spacetree approach described earlier.

The penalty coupling terms were integrated over a surface triangulation, automatically generated using marching cubes [68]. The implementation is a three-dimensional extension of the high-order parametrization algorithm used in Section 4.1, where a grid of $n_g \times n_g \times n_g$ points per sub-cell was used for the surface recovery. An example of the resulting triangulation is shown in Figure 18. Note that the surface mesh is irregular and that the triangles have severe aspect ratios. The high-order penalty terms are then integrated over the triangles by collapsing one side of the integration domain (bi-unit square) and mapping it to each triangle. As demonstrated by the two-dimensional ellipse benchmark the linear parametrization needs a fine resolution to achieve accurate results.

The numerical approximation for the axial stresses σ_{xx} is depicted in Figure 19. Note that half of the domain is removed to visualize the internal solution. For the unrefined meshes, the approximated stresses exhibit oscillations and non-physical jumps at the element boundaries. If the cube's mesh, which is affected by the concave geometry, is refined towards the two points, the stress concentration can be localized quickly. Similar to the two-dimensional case, the solution within the ellipsoid remains smooth, and does not require local mesh refinement.

The results show that the improvement in accuracy brought on by the multi-level *hp*-refinement scheme also carries over to three-dimensional problems. The example also demonstrates that a surface mesh for the interface discretization is easily obtained for smooth implicit geometries. The surface mesh is only used to integrate the penalty terms, and hence can be irregular or have severe aspect ratios, unlike what is required of an analysis-suitable finite element mesh.

4.4 Cylinder with cubical inclusion

The last numerical benchmark for us to consider is a three-dimensional problem with a vertex-edge singularity. We study a cylinder under axial tension, which has an embedded cubical inclusion. Due to the symmetry of the problem, we consider only one-eighth of the system. An overkill solution was generated for reference, using seven conforming hexahedral *p*-FEM elements with

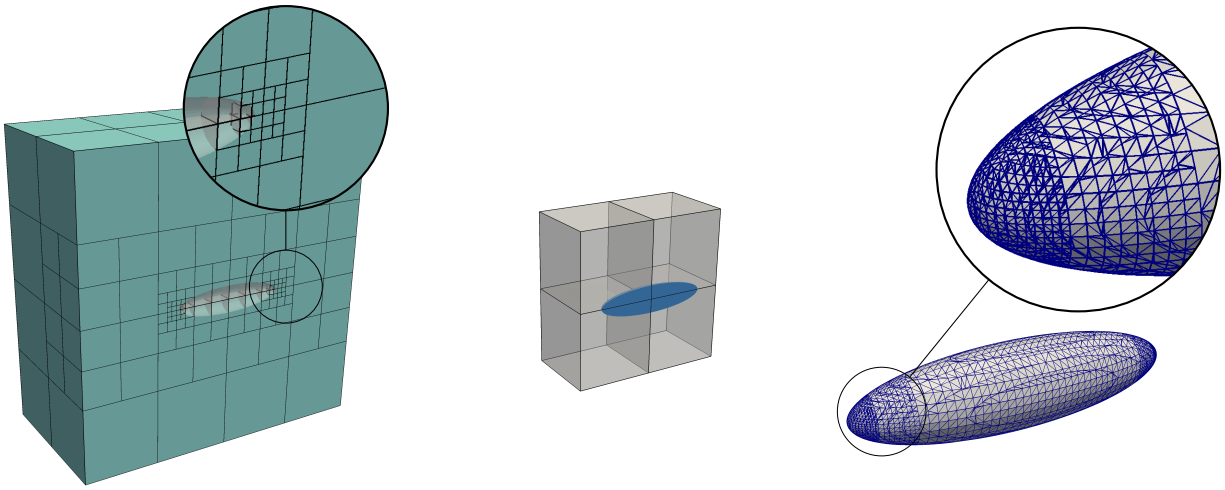

(a) Cube (b) Ellipsoidal inclusion (c) Coupling surface

Figure 18: Discretization of cube and ellipsoid using $k = 4$ - Section view

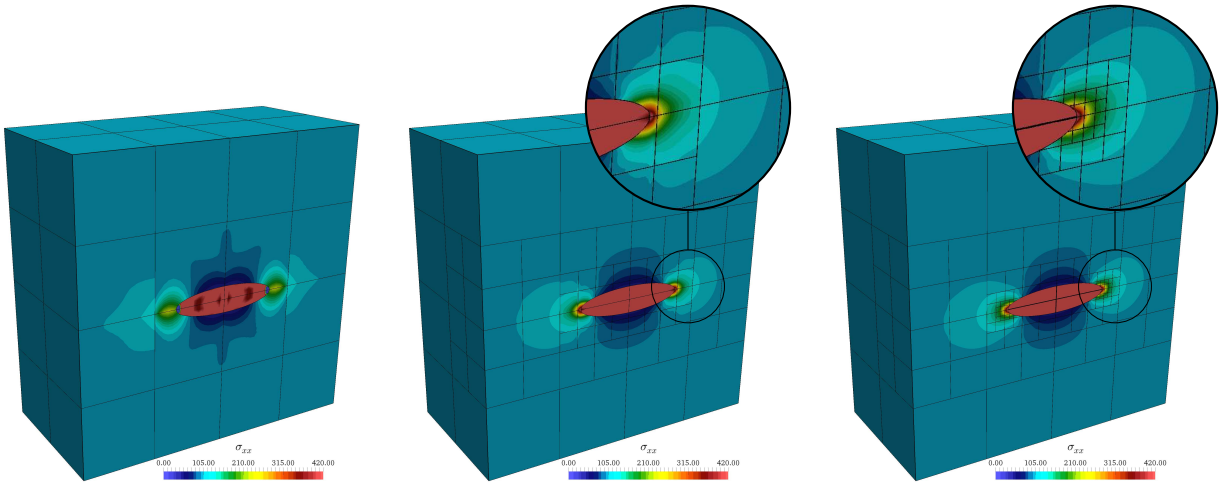

(a) $p = 8$, no refinement (b) $p = 6, k = 2$ (c) $p = 6, k = 4$

Figure 19: Cube with ellipsoidal inclusion: numerical approximation of axial stresses

exact blending, and a local refinement using the multi-level *hp*-scheme. Using 5 levels of refinement and a polynomial degree $p = 9$, the overkill discretization has 506,199 degrees of freedom. The reference strain energy for the setup shown in Figure 20a is

$$\mathcal{U}_{\text{ex}} = 1.037455 \times 10^3. \tag{23}$$

Two non-geometry-conforming FCM meshes were used for the cylinder and the cubical inclusion. A coarse base discretization of $2 \times 2 \times 2$ finite cells was used. The dimensions of the mesh were setup in such a way that the edges of the FCM meshes do *not* exactly coincide with the cube's edges. Similar to the two-dimensional benchmark presented in Section 4.2, the vertex-edge singularity affects the solution in both domains. Accordingly, both meshes were refined towards the singular edges, as shown in Figure 21. The symmetry boundary conditions were enforced in a classical manner on the corresponding faces of the FCM meshes.

The numerical approximation for the von-Mises stresses, σ_{eq}, is shown in Figure 22 for the

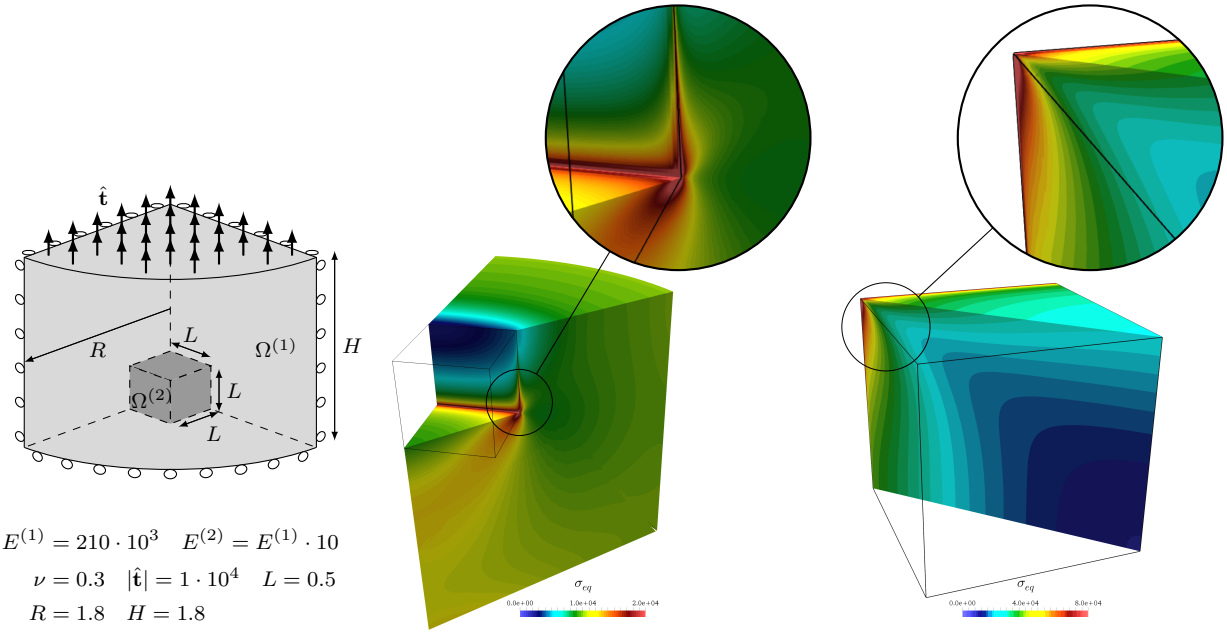

(a) Benchmark setup (b) Reference σ_{eq} in cylinder (c) Reference σ_{eq} in inclusion

Figure 20: Cylinder with cubical inclusion

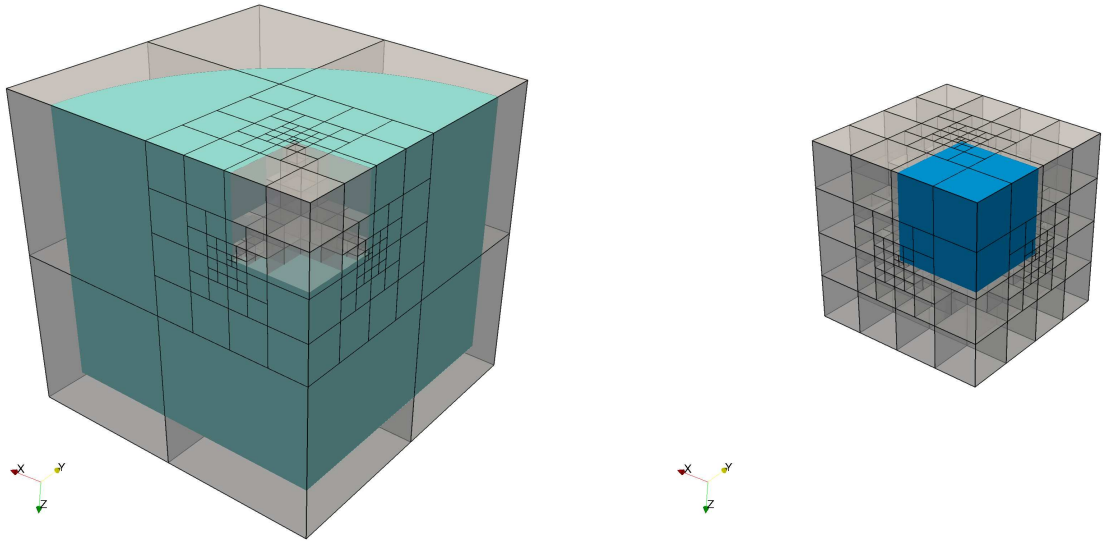

(a) Cylinder (b) Cubical inclusion

Figure 21: Discretization of cylinder and cubical inclusion, $k = 5$

cylinder and in Figure 23 for the cubical inclusion. The complete domain is clipped diagonally to reveal the solution at the vertex-edge singularity. Note that the geometry of the cylinder is rotated to show the solution along the singular edges. Whereas the solution obtained without refinement is highly oscillatory, the local refinement with the multi-level *hp*-scheme is able to confine the error to the finest level of sub-cells.

To assess the convergence behavior, a *p*-elevation study was carried out. The results, shown in Figure 24, indicate a slow algebraic rate for coupled FCM without refinement. Applying a local refinement starting with $k > 2$, a pre-asymptotic range can be identified as having a steeper

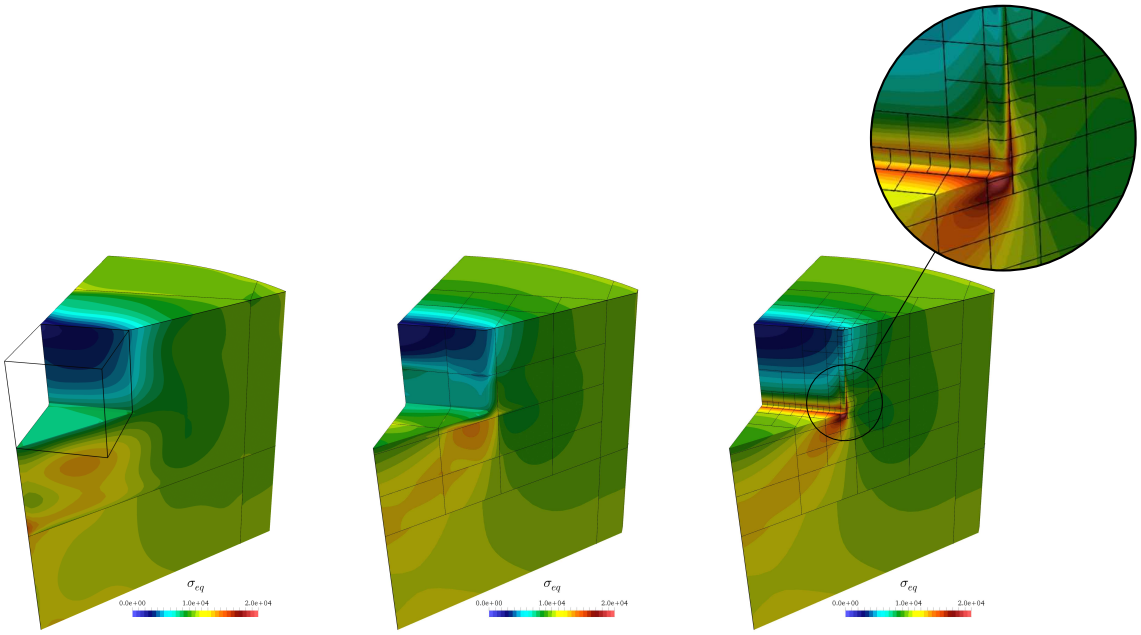

(a) $p = 8$, no refinement (b) $p = 4, k = 2$ (c) $p = 4, k = 5$

Figure 22: Cylinder: numerical approximation of σ_{eq}

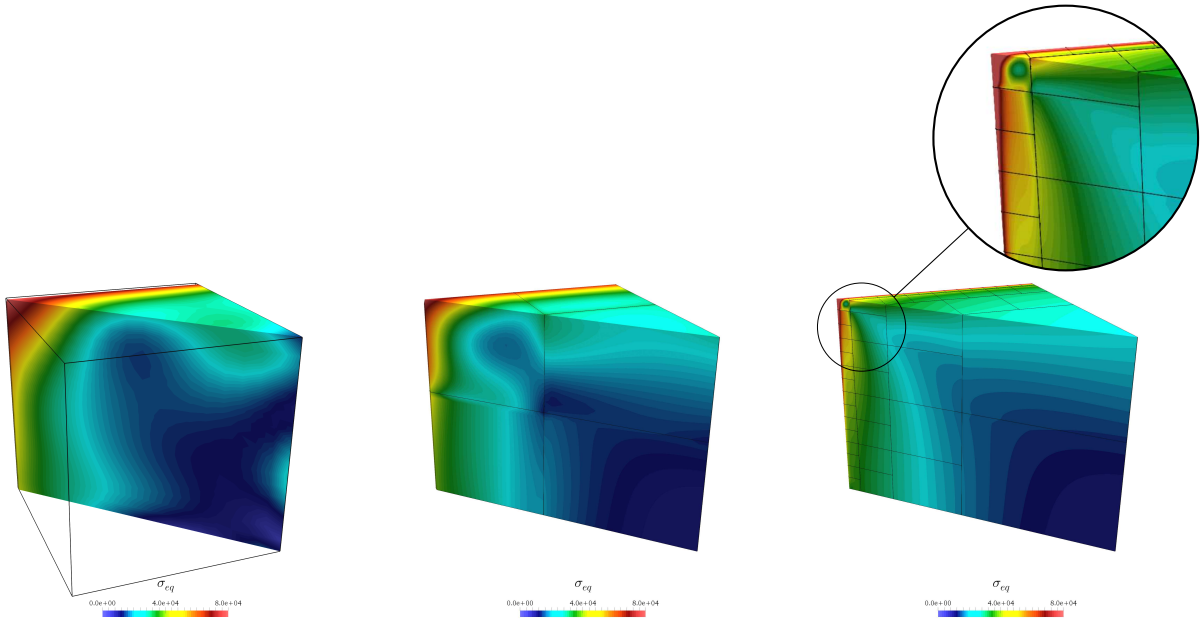

(a) $p = 8$, no refinement (b) $p = 4, k = 2$ (c) $p = 4, k = 5$

Figure 23: Cubical inclusion: numerical approximation of σ_{eq}

convergence rate. This allows the discretization with $k = 5$ to achieve an engineering accuracy of 1% error using a moderate polynomial order. This would require at least an order of magnitude more degrees of freedom with uniform h- or p-refinement. The pre-asymptotic convergence is also characterized as algebraic, albeit with a higher rate. For vertex-edge singularities, an anisotropic refinement scheme is necessary to attain exponential convergence [73], which was not applied in this work. However, this is not always feasible for the general case of embedded vertex-edge

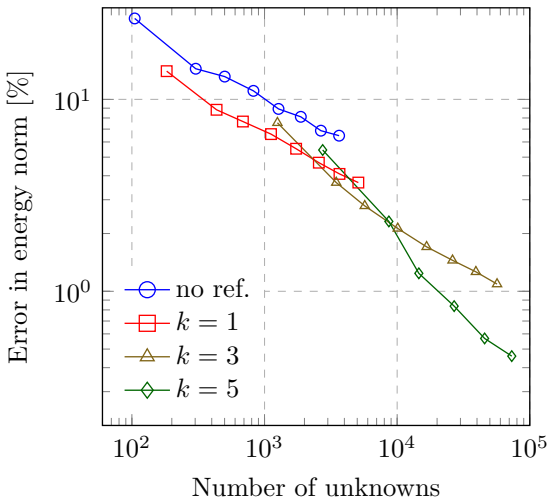

Figure 24: Cylinder with cubical inclusion: p-convergence for multi-level hp-refinement

singularities, where the (possibly curved) singular edges might not be parallel to edges of the mesh. Nevertheless, the improved algebraic convergence brought on by the multi-level hp-refinement allows for a significant reduction of the approximation error.

5 Application in biomechanics: vertebra with pedicle screws

Having investigated the convergence properties of the proposed refinement scheme for benchmark problems with stress concentration, we demonstrate the applicability of the refinement scheme for large-scale problems with complex three-dimensional geometries.

5.1 Problem setup

We consider a lower thoracic vertebra into which two cannulated pedicle screws are inserted. This is done as part of a procedure to stabilize a spinal segment of by fusing two vertebrae. We assume that the vertebra-screw interface is in full body contact, and model the contact as a material interface. We use the FCM with the weak coupling approach to solve the interface problem, and employ multi-level hp-refinement to resolve the high gradients around the screws' threads.

The geometry of the vertebra originates from a CT scan with high spatial resolution, carried out for a formalin-fixated specimen of the eleventh thoracic vertebra of an 84-year-old donor. She had dedicated her body for educational and research purposes to the local Institute of Anatomy prior to death, in compliance with the local institutional and legislative requirements. Images were acquired by using a whole-body 256-row CT scanner (iCT, Philips Medical Care, Best, The Netherlands) after 24 hours of degassing. Scan parameters were a tube voltage of 120 kVp, a tube load of 585 mAs, an image matrix of 1024×1024 pixels, and a field of view of 150 mm. Transverse sections were reconstructed with an interpolated voxel size of $146 \times 146 \times 146$ μm^3, and the intensity values of the CT images were calibrated with a reference phantom (Mindways Osteoporosis Phantom, San Francisco, CA, USA) to derive calcium hydroxyapatite values in (mg/cm^3).

The geometric model of the screws is a B-Rep CAD model of a Viper2-Screw (DePuy Synthes, Umkirch, Germany), provided by the manufacturer.

The material properties for the vertebra were resolved on the voxel level, as a threshold was set to the intensity values to distinguish the trabecular structure. The voxels that are identified as inside the bone were assigned a Young's modulus $E^{(1)} = 10$ GPa, and Poisson ratio $\nu = 0.3$, which are commonly used parameters in HR-pQCT based voxel FEM [75]. The material of the screws is titanium with $E^{(2)} = 100$ GPa and $\nu = 0.3$.

The considered load case simulates a pull-out test, where the vertebral body is clamped and the

screws are pulled axially outwards—as depicted in Figure 25. The Dirichlet boundary conditions are defined as clamping of superior and inferior end-plates, whereas the Neumann boundary condition is defined on the screws head, as a uniformly distributed axial force.

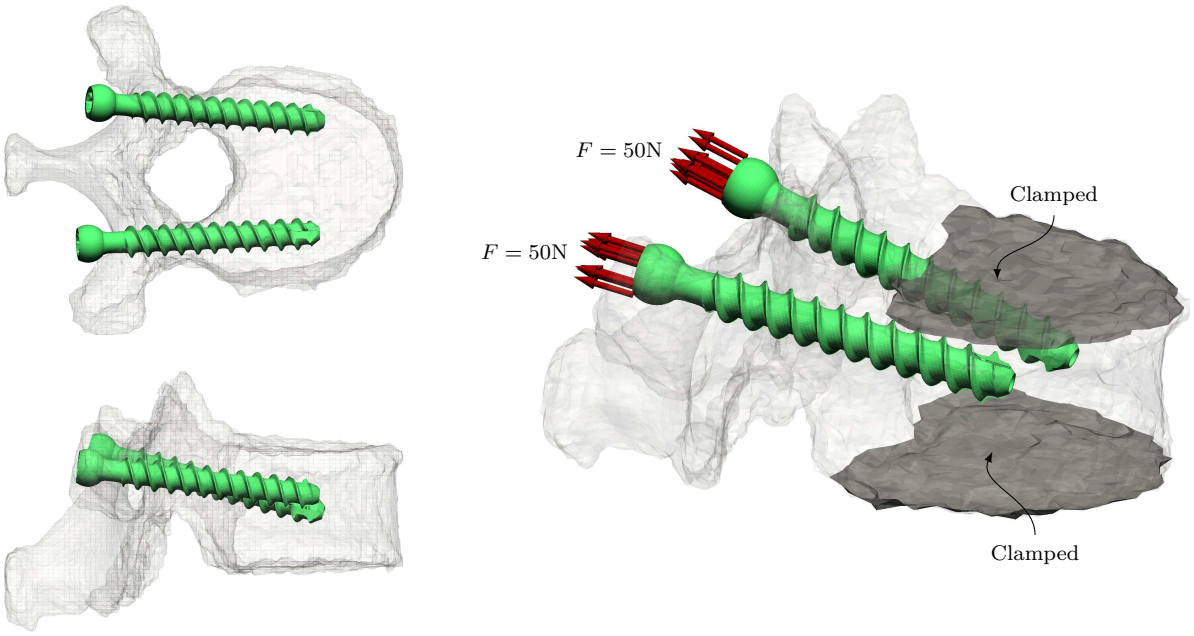

(a) Geometry (b) Boundary conditions

Figure 25: Pull-out test

5.2 Discretization

5.2.1 *hp*-FCM

The complete model of the vertebra-screw system was setup virtually by combining both geometric models as described in the following. First, the CT-scan was segmented using `ITK-SNAP` [76] to obtain a separate model of the vertebral body without the surrounding tissue. This segmentation only considers the outer boundary of the vertebra, and does not distinguish between cortical and trabecular bone.

The CAD model of the screws was suitably positioned in the same coordinate space. To carry out the inside-outside test on the CAD model, a watertight surface triangulation of the screws was created using `Rhinoceros` [77]. The surface triangulation was then utilized by a ray intersection algorithm [78], to determine whether a given integration point lies within the screws' geometry.

The first FCM mesh is used to discretize the embedding domain of the vertebral body. The mesh is axis-aligned with the CT-scan. The FCM model uses the geometric description based on the voxel data, as presented in [11, 49]. The geometry is defined by applying a threshold to the CT-scan to define the bone-structure and subtracting the screws through a Boolean operation. To account for the canals within the screws, an additional cylinder at the central axis of each screw was subtracted. The numerical evaluation of the domain integrals performs an intersection of the voxel grid with the refined FCM mesh to give the integration sub-cells.

To resolve the stress concentration, the mesh is refined using the multi-level *hp*-scheme towards the cusps of the screws' threads (concave geometry) and towards the re-entrant corner at the pedicle. The second FCM mesh is used to discretize the embedding domain of the screws. This mesh is refined towards the roots of the screws' threads, where stress concentration within the screws is expected. The spacetree scheme is used for the numerical integration on the second mesh.

The complete discretization results in a total of $1,322$ finite cells for first mesh and 155 finite cells for the second mesh. Using polynomial degree $p = 3$ and $k = 3$ refinement levels, the discretization has approximately 1.13 million degrees of freedom. The final refined meshes are depicted in Figure 26.

To control the ill-conditioning of the stiffness matrices, the fictitious domain in both meshes is given a stiffness of $E_{\text{fict}} = 10^{-4}$ GPa and $\nu_{\text{fict}} = 0.3$. The FCM model contains badly cut cells, which are expected in cases involving complex geometry, resulting in an ill-conditioned matrix even with relatively low polynomial degrees [79]. Hence, the chosen fictitious stiffness is relatively high. Nevertheless, as discussed at the end of Section 4.2, the introduced modeling error is expected to be much lower than the approximation error for this complex model.

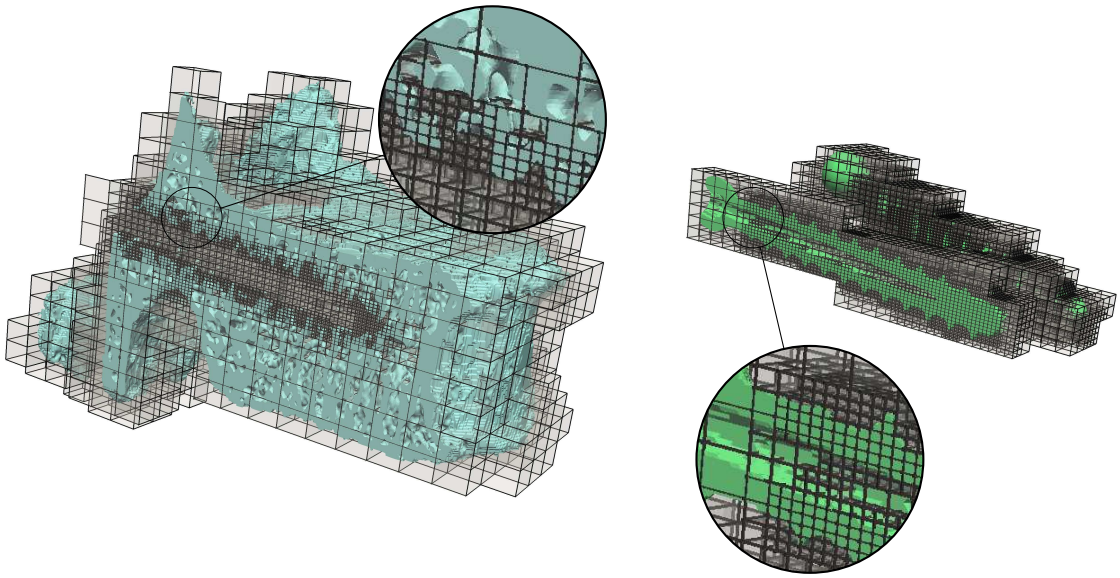

(a) Vertebra (b) Screws

Figure 26: Discretization using FCM

5.2.2 Material interface

The penalty approach was used to couple the meshes with $\beta_{12} = 10^9$. To integrate the coupling terms, a fine surface triangulations for the bone-screw interfaces was created using `Rhinoceros`, by meshing the outer surfaces of the screws. For an accurate and efficient evaluation of the surface integral, the triangulation was intersected with the refined finite cell grid, as shown in Figure 27. The triangles were then associated to the finite cell pairs form the two meshes, such that the coupling terms were added to the assembled stiffness matrices only once per finite cell pair.

5.2.3 Boundary conditions

Dirichlet boundary conditions were applied at the superior and inferior end-plates of the vertebra. These were weakly enforced using the penalty method, with the penalty terms integrated over a surface triangulation. The surface description was generated using the marching cubes algorithm and the segmentation. The penalty value was selected empirically, $\beta_D = 10^7$. The Neumann boundary conditions modeling the axial load on the screws were applied on a surface mesh created from the CAD model. The surface mesh was also intersected with the faces of the refined FCM mesh, to allow for a more accurate evaluation of the Neumann integral.

5.3 Solution and numerical results

The resulting system of linear equations was solved using the parallel direct solver `Intel`® `Pardiso` which is provided as part of the Intel Math Kernel Library [80]. The simulation was run on two

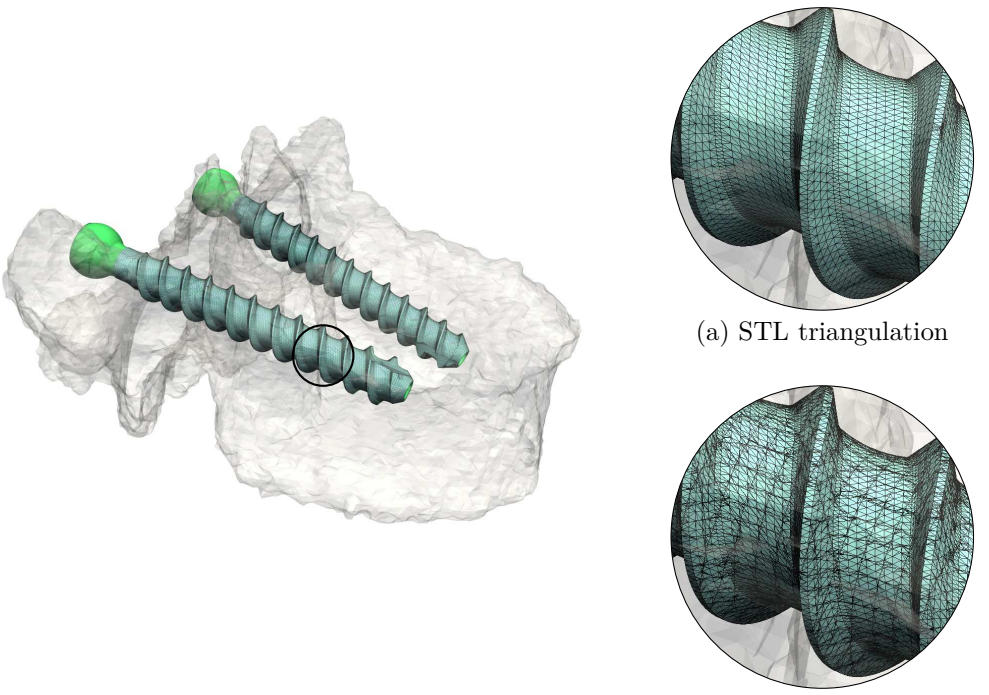

(a) STL triangulation

(b) Intersected triangulation

Figure 27: Discretization of coupling surface

eight-core Intel® Xeon® E5-2690 @ 2.9 GHz CPUs. The model takes approximately 40 minutes for the integration, solution and post-processing. Visualization was carried out using `ParaView` [81] on the same hardware.

The surface of the trabecular bone was recovered using `ParaView`, for an axial and a sagittal section. The recovered surface was used for the plots shown here. The computed displacements are depicted in Figure 28. It can be observed that the applied boundary conditions are fulfilled, and that the displacement field appears continuous across the interface.

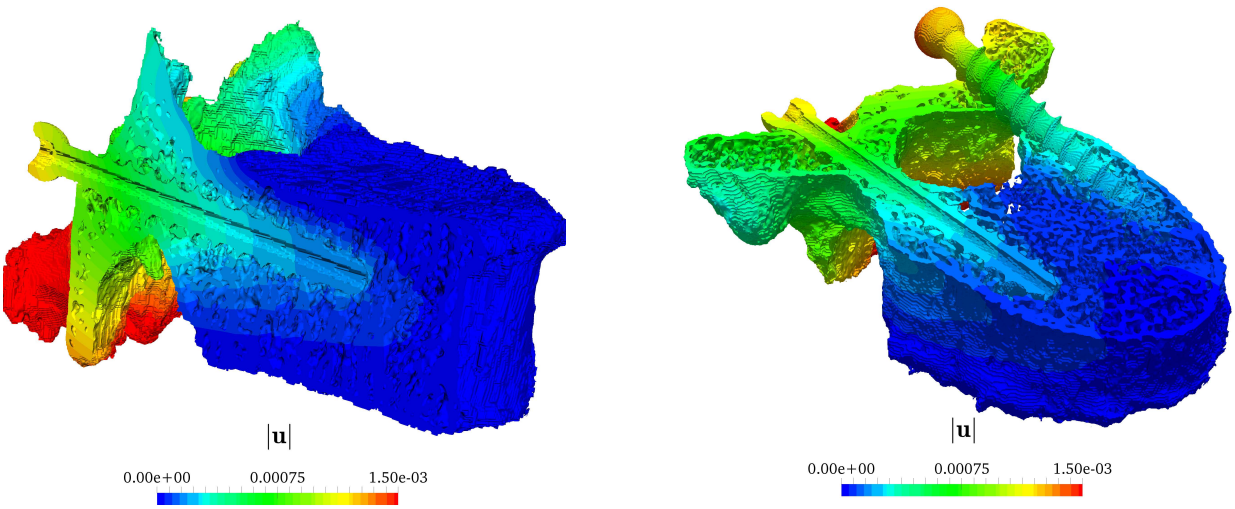

(a) Sagittal section

(b) Axial section

Figure 28: Vertebra with pedicle screws: numerical approximation of displacements

The von-Mises stresses calculated for the screws and the bones are depicted in Figure 29 for a sagittal and an axial section. The numerically approximated stresses are continuous within

the screw's geometry, with no discernible jumps at element boundaries or large scale oscillations. This indicates a low discretization error. High stresses are visible in the neck region between the applied load and the vertebra. Additionally, a bending stress distribution is observed for the helical thread, with the maximum stress occurring at the roots of the thread. In the vertebra, the stress concentration is localized around the screws, and at the pedicle's reentrant corner. Similar stress distribution patterns were also reported for simulations using micro-FE models [82–84].

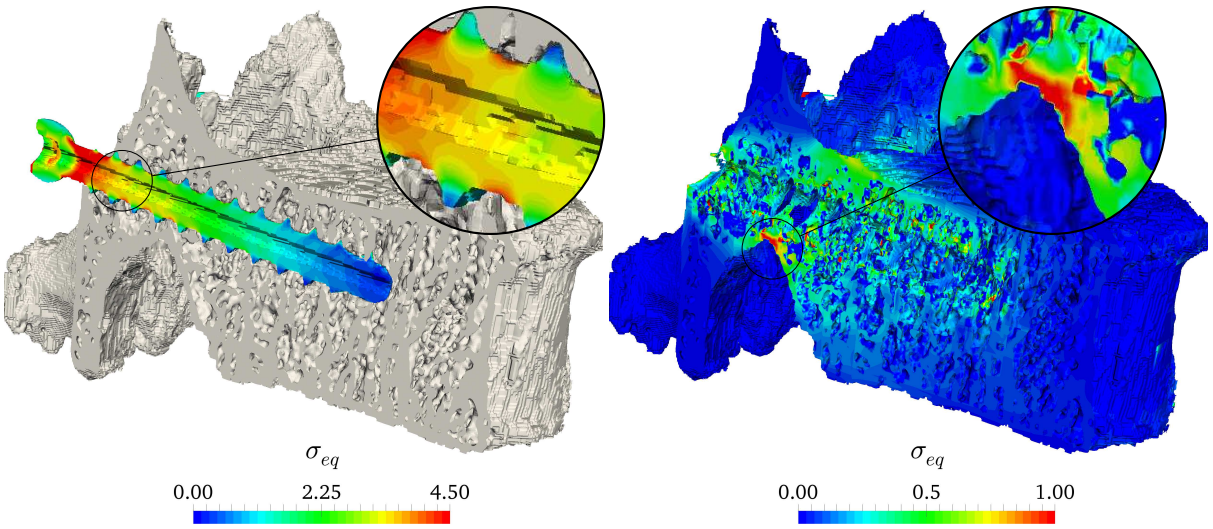

(a) Sagittal section

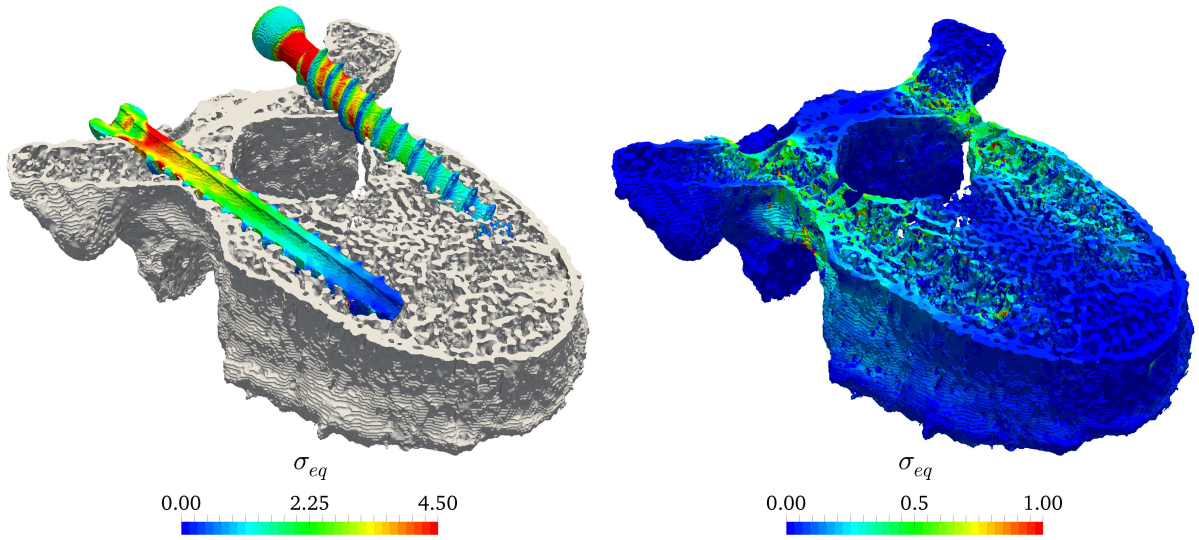

(b) Axial section

Figure 29: Vertebra with pedicle screws: numerical approximation of von-Mises stresses

Altogether, the numerical results are mechanically plausible, and correspond well to the applied load case. This demonstrates that the proposed refinement technique is suitable for solving problems involving complex three-dimensional geometries without the need to generate a boundary-conforming mesh. Furthermore, this application demonstrates the flexibility of the FCM in dealing with different types of geometric descriptions (CAD and image-based models), without the additional effort of re-parametrization.

6 Summary and Outlook

In the present paper, we demonstrated that the use of multi-level *hp*-adaptive refinement for the finite cell method in conjunction with a weak enforcement of the interface conditions constitutes an efficient and robust approach for the solution of material interface problems involving complex geometries. The proposed combination of discretization techniques circumvents volumetric mesh generation for domains with multiple materials, while maintaining—locally—a high model resolution.

Several numerical benchmarks of two- and three-dimensional material interface problems show that the convergence behavior is significantly improved by the multi-level *hp*-refinement scheme for problems with stress concentration and singularities. The results also highlight the importance of a high-order accurate geometric description of the interface and the geometric boundaries.

In the context of biomechanical problems, we demonstrated that the proposed scheme is directly applicable for the simulation of geometrically complex implant-vertebra models. Furthermore, we demonstrated that the flexible nature of the FCM in dealing with different types of geometric description allows for an easy combination of the image-based model of the vertebra and the CAD-based model of the screws. Additionally, the accurate approximation of the mechanical stresses, which was corroborated by several numerical benchmarks, further illustrates the potential of the proposed scheme as a simulation tool for bone-implant models.

The results presented here offer several possibilities for further developments. Future work will address the experimental validation of the vertebra model. Furthermore, extending the vertebra-implant model to include several vertebral bodies would allow for more realistic boundary conditions and loading scenarios. Also, the development of a robust and efficient error estimator for the FCM would open the door for automatic refinement.

Acknowledgements

The authors gratefully acknowledge the financial support of the German Research Foundation (DFG) under Grant RA 624/26-1, and of the European Research Council under Grant ERC-2014-StG 637164.

References

[1] T. J. R. Hughes, *The Finite Element Method: Linear Static and Dynamic Finite Element Analysis*. Mineola, NY: Dover Publications, 2000.

[2] K. J. Bathe, *Finite Element Procedures*. New Jersey: Prentice Hall, 2007.

[3] J. Parvizian, A. Düster, and E. Rank, "Finite cell method," *Computational Mechanics*, vol. 41, pp. 121–133, Apr. 2007.

[4] 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*, vol. 197, pp. 3768–3782, Aug. 2008.

[5] N. Zander, S. Kollmannsberger, M. Ruess, Z. Yosibash, and E. Rank, "The Finite Cell Method for linear thermoelasticity," *Computers & Mathematics with Applications*, vol. 64, pp. 3527–3541, Dec. 2012.

[6] D. Schillinger, M. Ruess, N. Zander, Y. Bazilevs, A. Düster, and E. Rank, "Small and large deformation analysis with the p- and B-spline versions of the Finite Cell Method," *Computational Mechanics*, vol. 50, pp. 445–478, Feb. 2012.

[7] S. Duczek, M. Joulaian, A. Düster, and U. Gabbert, "Numerical analysis of Lamb waves using the finite and spectral cell methods," *International Journal for Numerical Methods in Engineering*, vol. 99, pp. 26–53, July 2014.

[8] M. Elhaddad, N. Zander, S. Kollmannsberger, A. Shadavakhsh, V. Nübel, and E. Rank, "Finite Cell Method: High-Order Structural Dynamics for Complex Geometries," *International Journal of Structural Stability and Dynamics*, vol. 15, p. 1540018, Apr. 2015.

[9] M. Joulaian, S. Duczek, U. Gabbert, and A. Düster, "Finite and spectral cell method for wave propagation in heterogeneous materials," *Computational Mechanics*, vol. 54, pp. 661–675, Apr. 2014.

[10] D. Giraldo and D. Restrepo, "The spectral cell method in nonlinear earthquake modeling," *Computational Mechanics*, pp. 1–21, Aug. 2017.

[11] M. Ruess, D. Tal, N. Trabelsi, Z. Yosibash, and E. Rank, "The finite cell method for bone simulations: Verification and validation," *Biomechanics and modeling in mechanobiology*, vol. 11, pp. 425–37, Mar. 2012.

[12] C. Verhoosel, G. 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*, vol. 284, pp. 138–164, Feb. 2015.

[13] N. Zander, T. Bog, M. Elhaddad, R. Espinoza, H. Hu, A. Joly, C. Wu, P. Zerbe, A. Düster, S. Kollmannsberger, J. Parvizian, M. Ruess, D. Schillinger, and E. Rank, "FCMLab: A finite cell research toolbox for MATLAB," *Advances in Engineering Software*, vol. 74, pp. 49–63, Aug. 2014.

[14] 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*, vol. 200, pp. 3358–3380, Nov. 2011.

[15] M. Joulaian and A. Düster, "Local enrichment of the finite cell method for problems with material interfaces," *Computational Mechanics*, vol. 52, pp. 741–762, Oct. 2013.

[16] J. M. Melenk and I. Babuška, "The partition of unity finite element method: Basic theory and applications," *Computer Methods in Applied Mechanics and Engineering*, vol. 139, pp. 289–314, Dec. 1996.

[17] T. Strouboulis, K. Copps, and I. Babuv ska, "The generalized finite element method," *Computer methods in applied mechanics and engineering*, vol. 190, no. 32, pp. 4081–4193, 2001.

[18] T.-P. Fries and T. Belytschko, "The extended/generalized finite element method: An overview of the method and its applications," *International Journal for Numerical Methods in Engineering*, vol. 84, pp. 253–304, Oct. 2010.

[19] N. Moës, M. Cloirec, P. Cartraud, and J. F. Remacle, "A computational approach to handle complex microstructure geometries," *Computer Methods in Applied Mechanics and Engineering*, vol. 192, pp. 3163–3177, July 2003.

[20] T.-P. Fries, "A corrected XFEM approximation without problems in blending elements," *International Journal for Numerical Methods in Engineering*, vol. 75, pp. 503–532, July 2008.

[21] G. Legrain, R. Allais, and P. Cartraud, "On the use of the extended finite element method with quadtree/octree meshes," *International Journal for Numerical Methods in Engineering*, vol. 86, pp. 717–743, May 2011.

[22] G. Legrain, N. Chevaugeon, and K. Dréau, "High order X-FEM and levelsets for complex microstructures: Uncoupling geometry and approximation," *Computer Methods in Applied Mechanics and Engineering*, vol. 241-244, pp. 172–189, Oct. 2012.

[23] S. Soghrati, A. M. Aragón, C. Armando Duarte, and P. H. Geubelle, "An interface-enriched generalized FEM for problems with discontinuous gradient fields," *International Journal for Numerical Methods in Engineering*, vol. 89, pp. 991–1008, Feb. 2012.

[24] S. Soghrati and P. H. Geubelle, "A 3D interface-enriched generalized finite element method for weakly discontinuous problems with complex internal geometries," *Computer Methods in Applied Mechanics and Engineering*, vol. 217–220, pp. 46–57, Apr. 2012.

[25] S. Soghrati and H. Ahmadian, "3D hierarchical interface-enriched finite element method: Implementation and applications," *Journal of Computational Physics*, vol. 299, pp. 45–55, Oct. 2015.

[26] M. Safdari, A. R. Najafi, N. R. Sottos, and P. H. Geubelle, "A NURBS-based interface-enriched generalized finite element method for problems with complex discontinuous gradient fields," *International Journal for Numerical Methods in Engineering*, vol. 101, pp. 950–964, Mar. 2015.

[27] M. Safdari, A. R. Najafi, N. R. Sottos, and P. H. Geubelle, "A NURBS-based generalized finite element scheme for 3D simulation of heterogeneous materials," *Journal of Computational Physics*, vol. 318, pp. 373–390, Aug. 2016.

[28] J. Nitsche, "Über ein Variationsprinzip zur Lösung von Dirichlet-Problemen bei Verwendung von Teilräumen, die keinen Randbedingungen unterworfen sind," *Abhandlungen aus dem Mathematischen Seminar der Universität Hamburg*, vol. 36, pp. 9–15, July 1971.

[29] A. Hansbo and P. Hansbo, "An unfitted finite element method, based on Nitsche's method, for elliptic interface problems," *Computer Methods in Applied Mechanics and Engineering*, vol. 191, pp. 5537–5552, Nov. 2002.

[30] J. Dolbow and I. Harari, "An efficient finite element method for embedded interface problems," *International Journal for Numerical Methods in Engineering*, vol. 78, pp. 229–252, Apr. 2009.

[31] C. Annavarapu, M. Hautefeuille, and J. E. Dolbow, "A robust Nitsche's formulation for interface problems," *Computer Methods in Applied Mechanics and Engineering*, vol. 225–228, pp. 44–54, June 2012.

[32] W. Jiang, C. Annavarapu, J. E. Dolbow, and I. Harari, "A robust Nitsche's formulation for interface problems with spline-based finite elements," *International Journal for Numerical Methods in Engineering*, vol. 104, pp. 676–696, Nov. 2015.

[33] D. Schillinger, I. Harari, M.-C. Hsu, D. Kamensky, S. K. F. Stoter, Y. Yu, and Y. Zhao, "The non-symmetric Nitsche method for the parameter-free imposition of weak boundary and coupling conditions in immersed finite elements," *Computer Methods in Applied Mechanics and Engineering*, vol. 309, pp. 625–652, Sept. 2016.

[34] M. Ruess, D. Schillinger, A. I. Özcan, and E. Rank, "Weak coupling for isogeometric analysis of non-matching and trimmed multi-patch geometries," *Computer Methods in Applied Mechanics and Engineering*, vol. 269, pp. 46–71, Feb. 2014.

[35] A. Stavrev, L. H. Nguyen, R. Shen, V. Varduhn, M. Behr, S. Elgeti, and D. Schillinger, "Geometrically accurate, efficient, and flexible quadrature techniques for the tetrahedral finite cell method," *Computer Methods in Applied Mechanics and Engineering*, vol. 310, 2016.

[36] S. Kollmannsberger, A. Özcan, J. Baiges, M. Ruess, E. Rank, and A. Reali, "Parameter-free, weak imposition of Dirichlet boundary conditions and coupling of trimmed and non-conforming patches," *International Journal for Numerical Methods in Engineering*, vol. 101, pp. 670–699, Mar. 2015.

[37] V. Varduhn, M.-C. Hsu, M. Ruess, and D. Schillinger, "The tetrahedral finite cell method: Higher-order immersogeometric analysis on adaptive non-boundary-fitted meshes," *International Journal for Numerical Methods in Engineering*, vol. 107, pp. 1054–1079, Jan. 2016.

[38] S. Duczek, F. Duvigneau, and U. Gabbert, "The finite cell method for tetrahedral meshes," *Finite Elements in Analysis and Design*, vol. 121, pp. 18–32, Nov. 2016.

[39] S. Duczek and U. Gabbert, "The finite cell method for polygonal meshes: Poly-FCM," *Computational Mechanics*,

pp. 1–32, June 2016.

[40] W. Gui and I. Babuška, "The h, p and h-p versions of the finite element method in 1 dimension Part I: The error analysis of the p-version," *Numerische Mathematik*, vol. 49, pp. 577–612, Nov. 1986.

[41] W. Gui and I. Babuška, "The h, p and h-p versions of the finite element method in 1 dimension Part II: The error analysis of the h-and h-p versions," *Numerische Mathematik*, vol. 49, pp. 613–657, Nov. 1986.

[42] W. Gui and I. Babuška, "The h, p and h-p versions of the finite element method in 1 dimension Part III: The Adaptive h-p Version," *Numerische Mathematik*, vol. 49, pp. 659–683, Nov. 1986.

[43] P. Šolín and J. Červený, "Automatic *hp*-adaptivity with arbitrary-level hanging nodes," Tech. Rep. Research Report No. 2006-07, The University of Texas at El Paso, Department of Mathematical Sciences, 2006.

[44] L. Demkowicz, W. Rachowicz, and P. Devloo, "A fully automatic *hp*-adaptivity," *Journal of Scientific Computing*, vol. 17, pp. 117–142, Dec. 2002.

[45] N. Zander, T. Bog, S. Kollmannsberger, D. Schillinger, and E. Rank, "Multi-level *hp*-adaptivity: High-order mesh adaptivity without the difficulties of constraining hanging nodes," *Computational Mechanics*, vol. 55, pp. 499–517, Feb. 2015.

[46] N. Zander, *Multi-Level hp-FEM: Dynamically Changing High-Order Mesh Refinement with Arbitrary Hanging Nodes*. Ph.D. thesis, Technische Universität München, Munich, 2017.

[47] N. Zander, T. Bog, M. Elhaddad, F. Frischmann, S. Kollmannsberger, and E. Rank, "The multi-level *hp*-method for three-dimensional problems: Dynamically changing high-order mesh refinement with arbitrary hanging nodes," *Computer Methods in Applied Mechanics and Engineering*, vol. 310, pp. 252–277, Oct. 2016.

[48] N. Zander, M. Ruess, T. Bog, S. Kollmannsberger, and E. Rank, "Multi-level *hp*-adaptivity for cohesive fracture modeling," *International Journal for Numerical Methods in Engineering*, vol. 109, 2017.

[49] Z. Yang, M. Ruess, S. Kollmannsberger, A. Düster, and E. Rank, "An efficient integration technique for the voxel-based finite cell method," *International Journal for Numerical Methods in Engineering*, vol. 91, pp. 457–471, Aug. 2012.

[50] L. Kudela, N. Zander, T. Bog, S. Kollmannsberger, and E. Rank, "Efficient and accurate numerical quadrature for immersed boundary methods," *Advanced Modeling and Simulation in Engineering Sciences*, vol. 2, pp. 1–22, June 2015.

[51] L. Kudela, N. Zander, S. Kollmannsberger, and E. Rank, "Smart octrees: Accurately integrating discontinuous functions in 3D," *Computer Methods in Applied Mechanics and Engineering*, vol. 306, pp. 406–426, July 2016.

[52] M. Joulaian, S. Hubrich, and A. Düster, "Numerical integration of discontinuities on arbitrary domains based on moment fitting," *Computational Mechanics*, pp. 1–21, Mar. 2016.

[53] A. Abedian, J. Parvizian, A. Düster, H. Khademyzadeh, and E. Rank, "Performance of Different Integration Schemes in Facing Discontinuities in the Finite Cell Method," *International Journal of Computational Methods*, vol. 10, p. 1350002, June 2013.

[54] V. Thiagarajan and V. Shapiro, "Adaptively weighted numerical integration over arbitrary domains," *Computers & Mathematics with Applications*, vol. 67, pp. 1682–1702, May 2014.

[55] C. A. Felippa, "Introduction to Finite Element Methods," lecture notes, Department of Aerospace Engineering Sciences, University of Colorado at Boulder, Boulder, Colorado, USA, 2013.

[56] M. Ruess, D. Schillinger, Y. Bazilevs, V. Varduhn, and E. Rank, "Weakly enforced essential boundary conditions for NURBS-embedded and trimmed NURBS geometries on the basis of the finite cell method," *International Journal for Numerical Methods in Engineering*, vol. 95, pp. 811–846, Sept. 2013.

[57] I. Babuv ska, "The Finite Element Method with Penalty," *Mathematics of Computation*, vol. 27, pp. 221–228, Apr. 1973.

[58] C. Vinci, *Application of Dirichlet Boundary Conditions in the Finite Cell Method*. Master's thesis, Lehrstuhl für Computation in Engineering, Technische Universität München, 2009.

[59] A. Apostolatos, R. Schmidt, R. Wüchner, and K.-U. Bletzinger, "A Nitsche-type formulation and comparison of the most common domain decomposition methods in isogeometric analysis: Domain Decomposition Methods in Isogeometric Analysis," *International Journal for Numerical Methods in Engineering*, vol. 97, pp. 473–504, Feb. 2014.

[60] C. D. Mote, "Global-local finite element," *International Journal for Numerical Methods in Engineering*, vol. 3, pp. 565–574, Oct. 1971.

[61] E. Rank, "Adaptive remeshing and h-p domain decomposition," *Computer Methods in Applied Mechanics and Engineering*, vol. 101, pp. 299–313, Dec. 1992.

[62] D. Schillinger, A. Düster, and E. Rank, "The *hp-d*-adaptive finite cell method for geometrically nonlinear problems of solid mechanics," *International Journal for Numerical Methods in Engineering*, vol. 89, no. 9, pp. 1171–1202, 2012.

[63] 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*, vol. 3, p. 33, Dec. 2016.

[64] G. Királyfalvi and B. A. Szabó, "Quasi-regional mapping for the p-version of the finite element method," *Finite elements in analysis and design*, vol. 27, no. 1, pp. 85–97, 1997.

[65] T.-P. Fries and S. Omerović, "Higher-order accurate integration of implicit geometries," *International Journal for Numerical Methods in Engineering*, vol. 106, pp. 323–371, Jan. 2015.

[66] T. P. Fries, S. Omerović, D. Schöllhammer, and J. Steidl, "Higher-order meshing of implicit geometries—Part I: Integration and interpolation in cut elements," *Computer Methods in Applied Mechanics and Engineering*, vol. 313, pp. 759–784, Jan. 2017.

[67] S. Omerović and T.-P. Fries, "Conformal higher-order remeshing schemes for implicitly defined interface problems," *International Journal for Numerical Methods in Engineering*, vol. 109, pp. 763–789, Feb. 2017.

[68] C. Maple, "Geometric design and space planning using the marching squares and marching cube algorithms," in *2003 International Conference on Geometric Modeling and Graphics, 2003. Proceedings*, pp. 90–95, IEEE, 2003.

[69] T. Bog, N. Zander, S. Kollmannsberger, and E. Rank, "Weak imposition of frictionless contact constraints on automatically recovered high-order, embedded interfaces using the finite cell method," *Computational Mechanics*, Aug. 2017.

[70] B. A. Szabó, A. Düster, and E. Rank, "The p-version of the finite element method," in *Encyclopedia of Computational Mechanics* (E. Stein, ed.), Chichester, West Sussex: John Wiley & Sons, 2004.

[71] Z. Yosibash, "Numerical analysis on singular solutions of the Poisson equation in two-dimensions," *Computational Mechanics*, vol. 20, pp. 320–330, Sept. 1997.

[72] R. B. Kellogg, "On the poisson equation with intersecting interfaces," *Applicable Analysis*, vol. 4, pp. 101–129, Jan. 1974.

[73] I. Babuška and B. Guo, "Approximation properties of the $h-p$ version of the finite element method," *Computer Methods in Applied Mechanics and Engineering*, vol. 133, pp. 319–346, July 1996.

[74] M. Dauge, A. Düster, and E. Rank, "Theoretical and Numerical Investigation of the Finite Cell Method," *Journal of Scientific Computing*, vol. 65, pp. 1039–1064, Mar. 2015.

[75] D. H. Pahr, E. Dall'Ara, P. Varga, and P. K. Zysset, "HR-pQCT-based homogenised finite element models provide quantitative predictions of experimental vertebral body stiffness and strength with the same accuracy as μFE models," *Computer Methods in Biomechanics and Biomedical Engineering*, vol. 15, pp. 711–720, July 2012.

[76] P. A. Yushkevich, J. Piven, H. C. Hazlett, R. G. Smith, S. Ho, J. C. Gee, and G. Gerig, "User-guided 3D active contour segmentation of anatomical structures: Significantly improved efficiency and reliability," *NeuroImage*, vol. 31, pp. 1116–1128, July 2006.

[77] Robert McNeel, "Rhinoceros 3D." http://www.rhino3d.com/.

[78] S. Bindick, M. Stiebler, and M. Krafczyk, "Fast kd-tree-based hierarchical radiosity for radiative heat transport problems," *International Journal for Numerical Methods in Engineering*, vol. 86, pp. 1082–1100, June 2011.

[79] 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*, vol. 316, pp. 297–327, Apr. 2017.

[80] Intel, "Intel Math Kernel Library." http://software.intel.com/en-us/intel-mkl.

[81] U. Ayachit, *The ParaView Guide: A Parallel Visualization Application*. USA: Kitware, Inc., 2015.

[82] A. J. Wirth, J. Goldhahn, C. Flaig, P. Arbenz, R. Müller, and G. H. van Lenthe, "Implant stability is affected by local bone microstructural quality," *Bone*, vol. 49, pp. 473–478, Sept. 2011.

[83] A. J. Wirth, R. Müller, and G. Harry van Lenthe, "The discrete nature of trabecular bone microarchitecture affects implant stability," *Journal of Biomechanics*, vol. 45, pp. 1060–1067, Apr. 2012.

[84] D. Ruffoni, A. J. Wirth, J. A. Steiner, I. H. Parkinson, R. Müller, and G. H. van Lenthe, "The different contributions of cortical and trabecular bone to implant anchorage in a human vertebra," *Bone*, vol. 50, pp. 733–738, Mar. 2012.