




Gecko

Design for *IGA*-type
discretization workflows

Gecko Technical Report 3

DC1 - Nicolò Antonelli



This project has received funding from the European Union's Horizon Europe research and innovation programme under grant agreement No 101073106
Call: HORIZON-MSCA-2021-DN-01

Funded by the
European Union

Executive summary

This report summarizes the progress of my PhD research on the development of advanced computational methods based on the Shifted Boundary Method (SBM) within the Isogeometric Analysis (IGA) framework, with a primary focus on Computational Fluid Dynamics (CFD) and related diffusion problems. The research addresses three main areas: high-order unfitted discretizations for elliptic problems, stabilized formulations for viscous flows, and the numerical treatment of non-Newtonian and viscoplastic fluids.

The proposed framework resolves major challenges associated with immersed and nonconforming geometries by shifting boundary and interface conditions to surrogate boundaries and by reconstructing solution fields through high-order Taylor expansions. This strategy eliminates the need for cut-cell integration, mitigates the small cut-cell problem, and ensures numerical stability, optimal convergence rates, and favorable conditioning of the resulting linear systems. The use of spline-based basis functions enables the systematic exploitation of high regularity and exact geometry representation.

For incompressible flows, the methodology is combined with Variational Multi-Scale (VMS) stabilization techniques to achieve robust velocity–pressure coupling on unfitted isogeometric meshes. This approach yields stable and accurate simulations of Stokes and viscous flow problems in the presence of complex immersed boundaries and locally refined discretizations. The framework is further extended to viscoplastic and non-Newtonian models, including Bingham-type fluids, through the integration of stabilized and regularized formulations. Particular attention is devoted to the interaction between solution regularity, boundary treatment, and nonlinear constitutive behavior, demonstrating the critical role of smooth spline spaces in achieving optimal accuracy and convergence.

Recent developments also include the extension of the SBM to multipatch and nonconforming isogeometric configurations through the Gap–Shifted Boundary Method, enabling high-order, penalty-free coupling between independently parameterized patches without introducing additional degrees of freedom. This extension provides a flexible mechanism for local refinement and complex geometric assembly while preserving robustness and scalability.

Overall, this research contributes to bridging design and analysis workflows by enabling efficient, accurate, and geometrically flexible simulations for industrially relevant CFD and multiphysics applications. Future work will focus on extending these methods to fully nonlinear flow regimes, moving and deforming geometries, three-dimensional multipatch assemblies, and coupled multiphysics problems.

List of abbreviations

<i>IGA</i>	<i>IsoGeometric Analysis</i>
<i>FEM</i>	<i>Finite Element Method</i>
<i>SBM</i>	<i>Shifted Boundary Method</i>
<i>VMS</i>	<i>Variational Multi-Scale</i>
<i>CFD</i>	<i>Computational Fluid Dynamics</i>

Introduction

This report presents recent advancements in computational methods aimed at enhancing the efficiency, robustness, and accuracy of fluid dynamics simulations, with a particular focus on unfitted and embedded discretizations based on the Shifted Boundary Method (SBM) within the Isogeometric Analysis (IGA) framework. The primary objective is to develop and assess high-order, penalty-free numerical techniques that streamline the pre-processing phase and enable rapid, reliable, and precise flow analysis for complex geometries and immersed objects. This is especially relevant in industrial sectors such as automotive, aeronautical, and mechanical engineering, where reducing geometry preparation time and improving solver performance remain major challenges.

A central motivation of this work is the need to perform accurate numerical simulations directly on CAD-derived models that may contain trimmed surfaces, geometric gaps, or nonconforming parametrizations. Traditional body-fitted approaches often require extensive reparametrization and meshing efforts, which hinder efficient design-to-analysis workflows. In contrast, the unfitted SBM framework avoids explicit mesh conformity by shifting boundary and interface conditions to surrogate boundaries and reconstructing the solution through high-order Taylor expansions. This strategy eliminates the need for cut-cell integration, alleviates the small cut-cell problem, and leads to well-conditioned linear systems that are amenable to efficient iterative solvers.

The research presented in this report covers both low- and high-order immersed and embedded methods applied to Computational Fluid Dynamics (CFD) and related diffusion problems. Particular emphasis is placed on the integration of high-regularity spline spaces provided by IGA with penalty-free Nitsche formulations and advanced stabilization techniques. This combination enables the systematic exploitation of smooth basis functions, exact geometry representation, and weak enforcement of boundary and interface conditions, resulting in optimal convergence rates and enhanced numerical robustness.

Beyond classical immersed boundary treatments, recent developments within this work include the extension of the SBM framework to nonconforming and multipatch isogeometric configurations through the Gap-Shifted Boundary Method. This approach enables the accurate and stable coupling of independently parameterized patches without introducing additional degrees of freedom, while preserving favorable conditioning. It further provides a flexible mechanism for local h- and p-refinement and for the treatment of complex assemblies commonly encountered in industrial CAD models.

The overarching goal of this research is to establish a unified and scalable framework for the analysis of CFD problems, ranging from scalar convection-diffusion equations to viscous incompressible Navier-Stokes and viscoplastic flow models. The proposed methods are designed to remain robust in the presence of complex immersed boundaries, nonmatching interfaces, and locally refined discretizations. Particular attention is devoted to the interaction between boundary treatment, solution regularity, stabilization, and nonlinear constitutive behavior.

All developments are implemented and validated within the open-source Kratos-Multiphysics platform, in alignment with the objectives of the GECKO project to integrate advanced numerical methodologies into industrial simulation environments. This ensures reproducibility, scalability, and direct applicability to real-world engineering problems.

This report is structured as follows:

SBM in IGA for Laplacian Problems.

This section presents the formulation and analysis of high-order SBM discretizations within the IGA framework for elliptic problems. Emphasis is placed on penalty-free boundary enforcement, optimal convergence, and conditioning properties, highlighting the advantages of spline-based unfitted methods for complex geometries.

SBM in IGA for Stokes Flow.

This section addresses the application of SBM-IGA to incompressible viscous flows. The focus is on velocity–pressure coupling, Variational Multi-Scale (VMS) stabilization, and robustness with respect to immersed boundaries and local refinement. Numerical results demonstrate stable and accurate performance in challenging geometric configurations.

SBM in IGA for Non-Newtonian Fluids.

This section extends the framework to viscoplastic and non-Newtonian flow models, with particular emphasis on Bingham-type fluids. Stabilized and regularized formulations are combined with high-order spline discretizations to accurately capture nonlinear rheological behavior and to analyze the influence of solution regularity on convergence.

Multipatch and Embedded Coupling via the Gap–SBM.

This section introduces the extension of the SBM framework to nonconforming multipatch configurations. The Gap–SBM methodology is presented as a robust and flexible tool for interface coupling, local refinement, and complex geometric assembly, preserving high-order accuracy and favorable conditioning.

Overall, this document contributes to the objectives of the GECKO project by demonstrating the feasibility and benefits of integrating high-order unfitted and embedded isogeometric methods into industrial CFD workflows. It establishes a foundation for future developments in immersed and multipatch analysis, supporting increasingly complex geometries, multiphysics couplings, and large-scale simulations within a unified design-to-analysis framework.

1. The Shifted Boundary Method in IGA

IGA has transformed computational mechanics by bridging the gap between Computer-Aided Design (CAD) and Computer-Aided Engineering (CAE). Originally introduced by Hughes et al. [1,2,3,4,5], IGA enables precise geometric representation and high continuity across element boundaries [6], making it an effective tool for simulating complex geometries [7,8,9]. The foundation of IGA lies in B-Spline and NURBS basis functions, which allow for smooth transitions and local refinement, enhancing the accuracy and robustness of simulations [10,11]. Despite its advantages, traditional boundary-fitted IGA methods face challenges with complex geometries, such as the need for watertight geometries and high computational costs in handling trimmed models [12,13]. Immersed boundary IGA methods, such as the Finite Cell Method (FCM) [14,15] and Isogeometric Boundary Representation Analysis (IBRA) [16,17,18,19,20], alleviate some of these challenges by using non-boundary-fitted discretizations. However, these approaches are often hindered by issues like the small cut-cell problem [21,22], which leads to poor computational performance and difficulty in solver convergence.

The SBM [23,24], initially developed within the FEM context, offers a promising solution to these challenges. SBM simplifies integration over complex geometries by shifting boundary conditions to a surrogate boundary and modifying boundary values using Taylor expansions. This approach avoids the small cut-cell problem, maintains optimal accuracy, and enables significant simplification in mesh generation and refinement. Previous applications of SBM in FEM have demonstrated its effectiveness in areas such as elasticity and incompressible fluid dynamics [25,26,27,28]. This study pioneers the integration of SBM with IGA, proposing an alternative to IBRA that eliminates the need for trimming techniques. By treating the boundary as unfitted, the proposed approach simplifies pre-processing and retains the accuracy and continuity advantages of IGA. The structure includes an introduction to the SBM framework, its application to the Poisson problem, and numerical validations.[31].

The Shifted Boundary Method (SBM) introduces key innovations.

- Surrogate boundary representation: the true boundary, denoted as Γ , is replaced by a surrogate boundary, which is constructed using the edges of a Cartesian background grid. This simplification avoids complications associated with small cut-cells in boundary-fitted methods (see Figure 1).
- Taylor expansion for boundary conditions: boundary conditions are imposed on the surrogate boundary using a Taylor expansion of arbitrary order m . This approach accurately approximates the conditions that would have been applied on the true boundary, preserving the desired level of accuracy.
- Weak imposition of boundary conditions: the shifted boundary operator is employed to weakly enforce Dirichlet and Neumann boundary conditions, ensuring numerical stability and compatibility with the unfitted nature of the method.

These features collectively enable SBM to achieve robust and accurate solutions without the computational overhead associated with traditional boundary-fitted approaches.




Figure 1: An example of a domain Ω with a ring-like shape treated with the SBM in a Cartesian grid. Blue and red solid lines denote the true and surrogate boundaries respectively. The computational domain is shaded in light yellow, while brown denotes the intersected cells that are not part of the SBM computation.

The Taylor expansion plays a fundamental role in the SBM for accurately shifting boundary conditions from the true boundary to the surrogate boundary. This process involves performing an m -th order Taylor expansion of the variable of interest, u , at the surrogate boundary. Assuming u is sufficiently smooth within the strip between Γ and the surrogate, the expansion incorporates directional derivatives along the vector $\mathbf{d}$, capturing the influence of the distance between the true and surrogate boundaries.

The m -th order Taylor expansion includes a remainder term that becomes negligible as the distance $\|\mathbf{d}\|$ approaches zero. Dirichlet conditions defined on Γ are extended using the boundary shift operator, enabling the surrogate boundary to approximate the true boundary conditions.

$$u(\mathbf{x}) = u(\tilde{\mathbf{x}} + \mathbf{d}(\tilde{\mathbf{x}})) = u(\tilde{\mathbf{x}}) + \sum_{i=1}^m \frac{D_{\mathbf{d}}^i u(\tilde{\mathbf{x}})}{i!} + (R^m(u, \mathbf{d}))(\tilde{\mathbf{x}})$$

Neglecting the remainder term simplifies the boundary condition approximation to an m -th order shifted expression, which is subsequently enforced weakly in the SBM framework.

$$S_{\mathbf{d}}^m u(\tilde{\mathbf{x}}) := u(\tilde{\mathbf{x}}) + \sum_{i=1}^m \frac{D_{\mathbf{d}}^i u(\tilde{\mathbf{x}})}{i!}$$

This approach ensures both computational efficiency and numerical accuracy while avoiding the complexities associated with direct boundary alignment [31].

$$S_{\mathbf{d}}^m u \approx u_D \quad \text{on } \tilde{\Gamma}_h$$

Moreover, we have taken advantage of two ideas recently published:

1. **Penalty-Free formulation:** A penalty-free weak formulation [29] eliminates the need for penalty calibration, which is traditionally required for Nitsche's method. The SBM variational form ensures numerical stability and accuracy without the overhead of parameter tuning.
2. **Optimal surrogate boundary selection:** the surrogate boundary minimizes the distance function between the surrogate and true boundaries [30]. A parameter λ determines the inclusion of cut elements in the surrogate domain, with $\lambda=0.5$ found to be optimal for minimizing error.




Figure 2: Selection of the surrogate boundary. Two-dimensional complex embedded holes treated using the SBM. The red solid line denotes the surrogate boundary corresponding to $\lambda = 0$ and the green one to $\lambda = 0.5$ (optimal surrogate boundary).

In Figure 3 we report the convergence studies of the two cases in Figure 2 when applying only Dirichlet boundary conditions.





Figure 3: Selection of the surrogate boundary. Convergence study of two cases curvilinear and diamond of Figure 2 respectively, considering $\lambda = 0$ (red solid line) and $\lambda = 0.5$ (blue solid line). The order of the basis functions is indicated by the following symbols: triangles $p = 1$, squares $p = 2$, and stars $p = 3$.

As previously mentioned, B-Spline basis functions are utilized to represent both the geometry and the solution fields. The implementation of the SBM avoids the use of trimmed knot spans, which are characteristic of the IBRA approach (refer to the left panel of Figure 4). The right panel of Figure 4 presents a comparison between the SBM and IBRA methodologies using linear, quadratic, and cubic B-Spline basis functions. The results demonstrate that both methods achieve very similar error norms and exhibit optimal convergence rates across all orders of the basis functions.



(a) IBRA Computational set-up.



(b) Convergence rates.

Figure 4: SBM vs IBRA. Left: the diamond shape applying the IBRA technique; the green dots denote the integration points, while the solid blue line highlights the true boundary. Right: the convergence analysis of the same case, using $p = 1$ (triangular symbols), $p = 2$ (square symbols), and $p = 3$ (star symbols) and comparing the IBRA approach (red solid line) and the SBM (blue solid line) with the optimal surrogate boundary.

For more details and examples we refer to the paper “The Shifted Boundary Method in Isogeometric Analysis” [31], where the authors also analyze the case of Neumann boundary conditions.

2. SBM-IGA for Stokes fluids

This section explores the application of the SBM within the IGA framework to solve Stokes flow problems. The Stokes equations describe the behavior of viscous, incompressible fluids under low Reynolds number conditions, where inertial forces are negligible compared to viscous forces. These equations involve coupling between the velocity field and the pressure field, which introduces additional numerical complexity compared to the Poisson problem discussed in Section 1.

The Stokes problem consists of solving the following equations:

$$\begin{aligned} -\nabla \cdot \boldsymbol{\tau} + \nabla p + \mathbf{f} &= 0 \quad \text{in } \Omega \\ \nabla \cdot \mathbf{u} &= 0 \quad \text{in } \Omega \end{aligned}$$

together with the boundary conditions. The weak form after integration by parts:

$$\begin{aligned} \int_{\Omega} \nabla^s \mathbf{v} : \mathbf{C} : \nabla^s \mathbf{u} \, d\Omega - \int_{\Omega} p \nabla \cdot \mathbf{v} \, d\Omega - \int_{\partial\Omega} (\mathbf{C} : \nabla^s \mathbf{u} \cdot \mathbf{n}) \cdot \mathbf{v} \, d\Gamma + \\ + \int_{\partial\Omega} p(\mathbf{v} \cdot \mathbf{n}) \, d\Gamma = \int_{\Omega} \mathbf{f} \cdot \mathbf{v} \, d\Omega, \\ \int_{\Omega} q \nabla \cdot \mathbf{u} \, d\Omega = 0, \end{aligned}$$

where $\mathbf{v}$ is the test function for the velocity vector field, and q is the test function for the pressure scalar field.

The Stokes problem is solved within a **Variational Multi-Scale (VMS)** [32-33] framework, which provides a stabilized formulation to handle the coupling of the velocity and pressure fields. In the VMS approach:

- The solution space is decomposed into resolved (coarse) and unresolved (fine) scales.
- Stabilization terms are introduced to account for the effects of unresolved scales on the resolved solution, enhancing numerical robustness and accuracy.

$$\begin{aligned} \mathbf{u} &= \mathbf{u}_h + \tilde{\mathbf{u}}, & \tilde{\mathbf{u}} &= \tau_m \mathbf{R}_m, \\ p &= p_h + \tilde{p}, & \tilde{p} &= \tau_c R_c, \end{aligned}$$

where

$$\begin{aligned} \mathbf{R}_m &= \mathbf{f} + \nabla \cdot \boldsymbol{\tau}(\mathbf{u}_h) - \nabla p_h, \\ R_c &= -\nabla \cdot \mathbf{u}_h. \end{aligned}$$

The VMS method, combined with B-Spline basis functions in the IGA framework, allows for arbitrary order discretizations of the velocity and pressure fields. Unlike the simpler Poisson

problem, the Stokes equations involve higher-order terms that do not simplify directly, making the implementation of VMS more intricate. The weak formulation now is

$$\begin{aligned} \int_{\Omega} \nabla^s \mathbf{v} : \mathbf{C} : \nabla^s (\mathbf{u}_h + \tilde{\mathbf{u}}) d\Omega - \int_{\Omega} (p_h + \tilde{p}) \nabla \cdot \mathbf{v} d\Omega + \\ - \int_{\partial\Omega} (\mathbf{C} : \nabla^s \mathbf{u} \cdot \mathbf{n}) \cdot \mathbf{v} d\Gamma + \int_{\partial\Omega} p(\mathbf{v} \cdot \mathbf{n}) d\Gamma = \int_{\Omega} \mathbf{f} \cdot \mathbf{v} d\Omega, \\ \int_{\Omega} q \nabla \cdot (\mathbf{u}_h + \tilde{\mathbf{u}}) d\Omega = 0. \end{aligned}$$

Therefore we have three additional terms with respect to a formulation without the sub-scale:

$$\begin{aligned} \text{TERM (A)} &= - \int_{\Omega} (\nabla \cdot \mathbf{v})(-\tau_c \nabla \cdot \mathbf{u}_h), \\ \text{TERM (B)} &= \int_{\Omega} \tau_m q \nabla \cdot (\mathbf{f} + \nabla \cdot \boldsymbol{\tau}(\mathbf{u}_h) - \nabla p_h), \\ \text{TERM (C)} &= \int_{\Omega} \nabla^s \mathbf{v} : \mathbf{C} : \nabla^s \tilde{\mathbf{u}}. \end{aligned}$$

The implementation of the SBM for Stokes flow in IGA follows a similar strategy as in Section 1:

- **B-Spline basis functions:** The velocity and pressure fields are approximated using B-Splines of arbitrary order, ensuring smooth and accurate solutions.
- **Shifted boundary conditions:** Dirichlet and Neumann boundary conditions are imposed weakly on the surrogate boundary using Taylor expansions.
- **Stabilization terms:** VMS stabilization terms are added to the variational formulation to address numerical instabilities arising from the coupling between velocity and pressure fields.

The combination of SBM and VMS within the IGA framework enables the efficient and accurate solution of Stokes flow problems.

In addition to Stokes fluids, we introduced the use of **non-Newtonian fluids**, specifically the Bingham fluid model. The Bingham constitutive law describes a material that behaves as a solid when the stress magnitude is below a certain yield stress and flows as a viscous fluid when the stress exceeds the yield stress. The stabilized formulation of the Bingham fluid incorporates an exponential factor that smooths the transition.

For small values of the regularization parameter, the behavior transitions smoothly between solid-like and fluid-like phases. Figure 5 shows the Stress vs Strain diagram.




Figure 5: Stress vs. Strain Rate for the Bingham Model with Various Regularization Parameters (m). The graph illustrates the stress response of a Bingham fluid as a function of strain rate ($\gamma_{\dot{}}$) for different values of the regularization parameter m .

We can show some preliminary results obtained with a manufactured solution for a Stokes fluid having a Bingham constitutive law. The manufactured solution is a smooth non-polynomial function in both the components of the velocity and in the pressure field. The convergence studies in the L2 norm of the error for velocity and pressure is shown in Figure 6 and Figure 7 ($m = 10$ and $m = 300$ respectively) for linear, quadratic, and cubic B-Splines basis functions.




Figure 6: Convergence Study for Velocity and Pressure Errors with Yield Stress $\tau_0 = 100$ and regularization parameter $m=10$. The figure shows the error convergence for the velocity (left) and pressure (right) fields as the mesh size h decreases. Results are presented for B-Spline basis functions of degree $p=1$ (blue), $p=2$ (red), and $p=3$ (cyan).




Figure 7: Convergence Study for Velocity and Pressure Errors with Yield Stress $\tau_0 = 100$ and regularization parameter $m=300$. The figure shows the error convergence for the velocity (left) and pressure (right) fields as the mesh size h decreases. Results are presented for B-Spline basis functions of degree $p=1$ (blue), $p=2$ (red), and $p=3$ (cyan).

We observed that the exact solution must be smooth, including its derivatives, to achieve optimal convergence when using B-Splines. Unlike classical FEM, where the solution is typically continuous between elements, the IGA framework with B-Splines ensures continuity throughout the domain. If the solution exhibits discontinuities, particularly under nonlinear constitutive laws such as the Bingham model, the method fails to achieve the expected optimal convergence rates. Therefore, when applying IGA to problems with nonlinear laws, it is essential to ensure the smoothness of the solution to fully leverage the advantages of B-Splines.

We have also validated our code by testing it on the well-known lid-driven cavity problem, as illustrated in Figure 8, adapted from [34]. This benchmark problem is widely recognized in computational fluid dynamics as a standard test for assessing the accuracy and robustness of numerical methods applied to fluid flow simulations. The lid-driven cavity problem involves a square domain where the top boundary (lid) moves at a constant velocity, inducing circulation within the fluid. The remaining boundaries are stationary and impose no-slip conditions. This configuration creates a primary vortex at the center of the domain.




Figure 8: Cavity problem. The flow in a 2D cavity filled with a Bingham fluid. The problem is set up following Mitsoulis and Zisis [34]. Defining a square domain $\Omega = (0, H) \times (0, H)$, we impose a horizontal velocity $u_x = 1$ m/s on the $y = H$.

Since no analytical solution is available for the lid-driven cavity problem, we have compared our results with the benchmark studies presented in [34] and [35]. This comparison focuses on two critical aspects of the flow behavior for different values of the yield stress: the extent of the yielded and unyielded regions within the domain, and the position of the eye of the main vortex that forms in the central-upper zone.

The yielded region, defined by areas where the fluid behaves as a viscous material, and the unyielded region, where the fluid behaves like a solid due to insufficient stress to induce flow, are key features of viscoplastic flows. Accurately capturing these regions provides insight into the robustness and precision of the numerical method (see Figure 9 and Figure 10).




Figure 9: Yielded region. Comparison of the yielded region in the cavity flow problem. The Bingham number is equal to 20.0 and $m = 300$. Left: results with B-Splines with $p=2$ and 50x50 knot spans. Right: results in FEM from [35] using local refinement.




Figure 10: Yielded region. Comparison of the yielded region in the cavity flow problem. The Bingham number is equal to 20.0 and $m = 300$. Left: results with B-Splines with $p=2$ and 100x100 knot spans. Right: results in FEM from [35] using local refinement.

Additionally, the location of the vortex eye serves as an important indicator of the method's ability to predict the flow dynamics accurately. The vortex position is influenced by both the yield stress and the imposed lid motion, making it a sensitive measure for evaluating the performance of our implementation (see Figure 11).




Figure 11: Eye vertex position. Comparison of the eye vertex y -coordinate in the cavity flow problem varying the Bingham number, $m = 300$. Left: results with B-Splines with $p=1$ & $p=2$. Right: results in FEM from [34] with 40×40 divisions.

Our results exhibit strong agreement with those in [34] and [35], confirming the capability of our Shifted Boundary Method (SBM) within the IsoGeometric Analysis (IGA) framework to handle complex boundary conditions and accurately simulate viscoplastic flow behavior in this classical benchmark problem.

3. Multipatch of nonconforming patches in Isogeometric Analysis

In practical industrial applications, complex geometries are rarely described by a single smooth parametrization. Instead, CAD models are typically composed of multiple patches, each defined on its own parametric domain and connected through interfaces that may exhibit geometric gaps, overlaps, or nonmatching discretizations [36]. While multipatch representations enable flexible geometric modeling and local refinement, they introduce significant challenges for numerical analysis, particularly in the context of high-order isogeometric discretizations, where ensuring watertightness and interface compatibility remains difficult.

Classical multipatch IGA approaches rely on conforming discretizations and strong continuity constraints at interfaces, requiring compatible knot vectors and matching parametrizations [37,38]. These requirements impose restrictive constraints on geometry generation and substantially increase preprocessing effort, especially for trimmed or nonconforming CAD models. Alternative weak coupling techniques, such as mortar methods or Nitsche-based formulations, relax some of these restrictions [39], but still require explicit interface integration and careful stabilization, particularly in the presence of geometric inconsistencies and small cut elements [37,40].

To overcome these limitations, this work extends the Shifted Boundary Method to the coupling of nonconforming isogeometric patches through the Gap–Shifted Boundary Method (Gap–SBM) [41]. This approach enables high-order, fully embedded coupling between independent patches with arbitrary differences in mesh size, polynomial degree, and parametrization, without introducing additional degrees of freedom. Interface conditions are weakly enforced on surrogate boundaries using high-order Taylor expansions [42], while avoiding the integration of small cut elements and the associated ill-conditioning. As a result, the Gap–SBM provides a robust and flexible framework for multipatch coupling in industrial IGA workflows involving complex and nonmatching geometries.

3.1 Gap–SBM Framework for Interface Coupling

The key idea of the Gap–SBM is to reinterpret patch interfaces as internal boundaries that may not coincide geometrically. Consider two adjacent patches whose physical boundaries are separated by a small gap or overlap. Instead of enforcing continuity on the true interface, surrogate interfaces are constructed on the discretization grids of each patch.

For each patch, an internal surrogate boundary is defined by selecting element faces close to the true interface. A closest-point projection operator maps points on the surrogate interface to their corresponding locations on the opposite patch. The distance vector between the true and surrogate interfaces is then used to reconstruct solution fields through high-order Taylor expansions, following the standard SBM philosophy.




Figure 12: Gap-SBM. The dashed black arrows symbolize the extension of the basis functions from the surrogate boundary into the gap region. Left: Extension of the basis functions from a surrogate point to integrate the gap volume. Right: Extension of the basis functions from a surrogate point to impose the boundary conditions on an approximation of the true boundary.

Interface conditions, such as continuity of primal variables and fluxes, are imposed weakly on the surrogate interfaces using shifted operators. This avoids direct integration over mismatched or disconnected geometries and eliminates the need for mesh conformity.

The resulting formulation preserves the high-order accuracy of IGA while enabling flexible coupling between independently parameterized patches.

3.2 Weak Enforcement of Interface Conditions

The coupling between adjacent patches is enforced through a penalty-free Nitsche-type formulation combined with the Gap-Shifted Boundary Method. This approach enables a fully embedded and high-order accurate treatment of interfaces that do not geometrically coincide, while avoiding the introduction of penalty parameters and preserving numerical stability.

For clarity, we focus on a representative configuration consisting of two patches: an inner patch containing a locally refined region and an outer patch surrounding it with an independent discretization. The two patches may differ in mesh size, polynomial degree, and parametrization. Although this section considers a two-patch configuration for illustration, the proposed methodology naturally extends to an arbitrary number of patches.

An illustrative example is shown in Figure 13, where the outer patch is discretized with a relatively coarse mesh, while the inner patch features a finer discretization and contains an embedded two-dimensional silhouette of the Stanford Bunny. It is emphasized that the presence of an embedded object is optional, and the focus of the proposed approach is on the coupling of nonconforming patches rather than on immersed boundary treatment.




Figure 13: Coupling using the Gap-SBM. In light-blue is the background patch (patch B) and light-yellow the refined inner patch (patch A).

In this framework, surrogate interfaces are constructed for both patches. Figure 13a shows the surrogate inner boundary (orange) and the surrogate outer boundary (green) associated with the refined inner patch. These surrogate boundaries provide the geometric support for the shifted interface operators. Figure 13b illustrates the volume integration points of both patches, together with the internal surrogate boundary of the coupling region (blue). This boundary represents the effective interface on which the weak coupling conditions are imposed.

To accurately integrate the gap region between the patches, the gap elements are subdivided into simple geometric entities. Figure 13c depicts the subdivision of the gap into Coons-type triangles and quadrilaterals (purple dashed lines), together with additional integration points (black dots) introduced to ensure continuity within the gap. These supplementary points are

required to correctly transfer information across the interface and to maintain consistency of the discretization.

Figure 13d shows the final volume integration points inside the gap region (red dots), which has been fully partitioned using Coons subdivisions. Within each sub-entity, high-order quadrature rules are constructed. No additional degrees of freedom are introduced for the integration of the gap. Instead, the basis functions of the outer patch are extended across the gap using Taylor expansions, in accordance with the SBM philosophy.

The approximate location of the coupling interface is defined using a convex hull construction of the embedded object or refined region. In the present example, the interface is obtained from an enlarged convex hull of the bunny silhouette and serves as the effective “true” boundary between the two patches. This choice reduces the number of degrees of freedom in the inner patch and facilitates strong local refinement. However, this is a design choice rather than a limitation of the method, as the coupling interface can be arbitrarily defined without affecting the validity of the formulation.

In the configuration shown in Figure 13, the interface is body-fitted with respect to the inner patch, while it is unfitted with respect to the outer patch. Other configurations are equally admissible, including fully unfitted or mixed fitted–unfitted couplings. The coupling region, highlighted in dark green in Figure 13, corresponds to the optimal surrogate boundary of the convex hull with respect to the inner patch.

After defining the interface, the inner patch admits a simple body-fitted coupling, whereas the outer patch is constructed independently. Its mesh size, polynomial degree, and parametrization are selected by the user, without imposing any compatibility constraints on knot vectors or control-point layouts. The only geometric requirement is that the coupling interface acts as the true internal boundary of the outer patch. From this boundary, the internal surrogate boundary of the outer patch is constructed, as shown in Figure 13b.

The theory of the Gap–SBM is then applied to accurately integrate the gap region and impose the coupling conditions. By looping over the surrogate boundary of the outer patch, the gap is iteratively partitioned into Coons-type elements until it is fully filled. Since the gap is always bounded by two surrogate boundaries, it remains geometrically simple, even in three-dimensional extensions of the method.

A key feature of this formulation is that the gap region can be integrated exactly. Unlike standard applications of the Gap–SBM to complex immersed geometries, no approximation of the coupling boundary is required in this setting. The interface itself is geometrically aligned with the surrogate boundaries, allowing the entire gap to be partitioned precisely and integrated without additional geometric reconstruction.

The weak coupling formulation resulting from this procedure preserves the high-order accuracy of the underlying isogeometric discretization and leads to symmetric and well-conditioned linear systems. Numerical experiments confirm that optimal convergence rates are recovered and that the coupling remains stable with respect to mesh refinement, interface misalignment, and gap thickness.

Finally, it is worth noting that this multipatch configuration enables flexible dynamic applications. For example, in CFD simulations involving rotating or moving components, the inner patch may be assigned to move with the object while the outer patch remains fixed. In this case, only the

coupling interface must be updated over time, avoiding re-meshing and enabling efficient simulation of fluid–fluid and fluid–structure interactions.

3.3 Convergence Results

We begin by assessing the accuracy of the proposed multipatch coupling in a controlled numerical setting. To this end, we consider a manufactured solution on a simple square domain and compare three discretization strategies: (a) a reference single-patch, body-fitted configuration, (b) a two-patch body-fitted configuration with matching interfaces, and (c) a two-patch embedded configuration using the Gap–SBM.




Figure 14: Mesh configurations used for the assessment of the multipatch coupling.

All configurations are solved using the penalty-free Nitsche interface formulation. A uniform step-by-step refinement of the underlying knot vectors is performed, and the error is evaluated in the L^2 -, L^∞ -, and H^1 -norms, as well as in terms of the L^2 -error versus the total number of degrees of freedom (DOFs). The objective of this study is to verify that the multipatch coupling preserves optimal convergence rates, to assess the impact of gap integration on accuracy, and to confirm the stability of the penalty-free formulation.

The mesh layouts for the three configurations are shown in Figure 14. The single-patch discretization is uniform and body-fitted (Figure 14a). The multipatch body-fitted case employs two adjacent patches with identical mesh sizes (Figure 14b). In contrast, the multipatch Gap–SBM configuration uses an independent parametrization for the inner patch, leading to a nonmatching interface and the presence of a strip of gap elements between the two patches (Figure 14c).

The corresponding convergence curves are reported in Figure 15. In addition to the three main configurations, an additional curve is included for the Gap–SBM case that accounts for the contribution of the gap elements. These results are highlighted in green.




Figure 15: Convergence study for the manufactured solution using the single-patch, the multipatch body-fitted coupling, and the multipatch Gap-SBM coupling. The Gap-SBM results including gap element contributions are highlighted in green.

The L^2 -error versus mesh size is shown in Figure 15a. All configurations exhibit optimal convergence rates of order $p + 1$ for polynomial degrees $p = 1, 2, 3$. The single-patch and multipatch body-fitted curves almost coincide across all refinement levels, indicating that the interface coupling does not introduce any measurable loss of accuracy. The multipatch Gap-SBM configuration also achieves optimal convergence, although a slight offset is observed, particularly for higher polynomial degrees.

A similar behavior is observed for the L^∞ -norm in Figure 15b. For the H^1 -norm, shown in Figure 15c, the multipatch configurations nearly coincide with the single-patch reference, demonstrating that the method accurately captures gradient information across nonmatching interfaces.

The error curve corresponding to the Gap-SBM configuration including the gap elements is slightly higher than the standard Gap-SBM curve, as expected, since the Taylor expansion of the basis functions is applied inside the gap region. Nevertheless, the two curves remain parallel, confirming that the approximation does not affect the asymptotic convergence behavior.

Figure 15d shows the L^2 -error plotted against the square root of the total number of DOFs. The single-patch configuration achieves the lowest error for a given number of DOFs, while the multipatch configurations exhibit comparable efficiency, with only a limited increase in DOFs due to interface contributions.

Overall, these results demonstrate that the proposed penalty-free multipatch coupling is robust, fully high-order accurate, and practically comparable in accuracy to a monolithic single-patch discretization. The Gap-SBM introduces only a modest and controlled increase in the error constant, while providing significant flexibility in terms of independent parametrization, local refinement, and geometric complexity.

3.5 Condition Number Analysis

To assess the numerical robustness of the proposed Gap–SBM coupling strategy, a condition number analysis is performed and compared with a standard single-patch configuration. The condition number of the stiffness matrix is a key indicator of numerical stability, as it quantifies the sensitivity of the linear system to perturbations. Large condition numbers are typically associated with poor solver performance and slow convergence of iterative methods.

In unfitted discretizations, small cut elements are known to cause severe ill-conditioning, since basis functions with very small active supports may lead to nearly linearly dependent system rows. The Shifted Boundary Method avoids this issue by evaluating basis functions on surrogate boundaries and by not enriching the approximation space inside cut elements. The same stabilization mechanism is inherited by the Gap–SBM formulation. Although integration is performed over the geometric gap between patches, no additional degrees of freedom are introduced, and all contributions are projected onto existing basis functions through Taylor expansions. As a result, the stiffness matrix is expected to retain the favorable conditioning properties of standard IGA discretizations.

Figure 16 reports the results of the condition number analysis for both single-patch and multipatch Gap–SBM configurations. The test setup consists of a square single-patch domain and a two-patch configuration with slightly mismatched mesh sizes, ensuring the presence of a gap and activating the Gap–SBM coupling mechanism.




Figure 16: Condition number and accuracy analysis for single-patch and Gap–SBM multipatch configurations of Fig. 5. The dashed colored lines indicate step-by-step convergence. Square, circle, and diamond markers denote polynomial degrees $p = 1$, $p = 2$, and $p = 3$, respectively.

Figure 16a shows the L^2 -error versus mesh size h , confirming that optimal convergence is preserved even for very fine discretizations. Figure 16b presents the condition number as a function of h . In all cases, the condition number exhibits the expected algebraic growth proportional to h^{-2} , with no indication of exponential blow-up. This behavior confirms the

absence of small cut-cell instabilities and indicates that all active basis functions remain sufficiently supported.

The step-by-step convergence results for the Gap–SBM configuration are also included in Figure 16b and show that the conditioning remains stable throughout the refinement process, with only minor oscillations that become slightly more pronounced for higher polynomial degrees.

Figure 16c reports the condition number versus the total number of degrees of freedom. The results follow the expected algebraic scaling with respect to the system size. For polynomial degrees $p = 2$ and $p = 3$, a plateau region is observed for coarse discretizations, where the condition number remains nearly constant before entering the asymptotic growth regime. In particular, for cubic basis functions, this plateau persists even for the finest mesh considered. Although the origin of this intermediate behavior is not yet fully understood, it further highlights the robustness of the proposed formulation when using high-order discretizations.

The increase in the absolute value of the condition number with the polynomial degree is not related to small cut-cell instabilities, which are effectively avoided by the SBM framework. Instead, this behavior is associated with the large support of B-Spline basis functions. Near embedded or surrogate boundaries, the effective support of the basis functions is truncated, so that only a fraction of their nominal support contributes to the stiffness matrix. While this fraction remains significant for low-order splines, it decreases with increasing polynomial degree, leading to larger condition numbers. This phenomenon can be interpreted as a small-support effect.

In practice, this effect does not lead to pathological ill-conditioning, but it may influence the efficiency of iterative solvers for higher-order discretizations ($p \geq 3$). Nevertheless, the overall results demonstrate that the Gap–SBM coupling preserves the favorable conditioning properties of standard IGA methods and provides a numerically robust framework for multipatch simulations.

4. 3D workflow in Kratos Multiphysics

This section describes the three-dimensional workflow implemented in the Kratos Multiphysics framework for the application of the Shifted Boundary Method within the Isogeometric Analysis setting. The objective is to outline the main steps required to construct and solve SBM-IGA formulations in 3D, with particular attention to geometry handling, surrogate boundary generation, discretization, and numerical integration within a practical computational environment. After presenting the workflow, some representative three-dimensional results are shown to demonstrate the performance of the proposed framework for both Laplacian problems and incompressible Navier–Stokes flows.

Within the Kratos Multiphysics framework, the three-dimensional workflow is designed to clearly separate the geometric description of the physical domain from the numerical discretization used for the solution. The fundamental idea is that the problem is not solved on a body-fitted mesh conforming to the true boundary, but rather on a regular IGA background discretization, while the physical boundary is represented in an immersed manner and replaced, at the discrete level, by a surrogate boundary aligned with the background grid. This strategy avoids the construction of complex conforming meshes while preserving an accurate geometric description and a consistent numerical formulation. The overall workflow can be interpreted as a sequence of geometric and numerical operations, starting from the import of the physical boundary and ending with the construction of the SBM volume and boundary entities used in the discrete problem.

Import of the boundary geometry

The starting point of the workflow is the definition of the physical boundary of the domain. In practice, this boundary may be provided either as a spline-based representation or as a triangulated surface, for instance derived from STL data. In the three-dimensional setting, the most natural input is therefore a closed surface delimiting the physical volume.

From a conceptual point of view, this geometry represents only the true boundary of the problem. It does not yet define the discretization on which the unknown field will be solved. Instead, it acts as the geometric reference required to identify which portions of the background domain are intersected by the physical boundary, to determine which regions belong to the interior or exterior of the physical domain, to establish the projection from the surrogate boundary to the true boundary, and to construct the SBM boundary conditions. At this stage, a consistent orientation of the boundary surface is essential, since the local orientation of the normals subsequently enters both the inside/outside classification and the construction of boundary terms.

Construction of the IGA background domain

Once the boundary description has been imported, a regular background domain is constructed, typically in the form of a NURBS volume fully enclosing the physical domain. This background volume is defined in terms of a physical bounding box, the corresponding parametric domain, the polynomial degree in the three parametric directions, and the number of knot spans along each direction.

The resulting structure is a regular three-dimensional grid of knot spans, which may be interpreted as a structured partition of the background volume into IGA cells. These cells are not fitted to the physical geometry; rather, the physical boundary may cut them in an arbitrary manner. This is a key aspect of the immersed formulation. The NURBS background volume therefore provides the geometric and functional support of the IGA discretization, the structured partition used to identify the regions affected by the immersed boundary, and the support on which the surrogate boundary is eventually constructed.

Detection of intersections between the boundary surface and the knot-span grid

The next stage consists of identifying how the true boundary intersects the background discretization. In the three-dimensional setting, this amounts to determining which cells of the background grid are cut by the physical surface.

Since the imported surface may be relatively coarse or may not be aligned with the scale of the knot-span partition, a local geometric refinement of the boundary representation is first introduced. The purpose of this refinement is not to improve the numerical approximation itself, but to make the detection of intersections with the background grid sufficiently robust. If a surface triangle or patch spans several knot spans or intersects them too coarsely, it is recursively subdivided until its size becomes compatible with the local scale of the background discretization.

As a result of this stage, a first geometric classification of the cells is obtained. Some cells are clearly intersected by the physical boundary, others are clearly far from the interface, while the remaining ones require further classification in terms of their actual inclusion in the physical domain.

Inside/outside classification of the background cells

Once the cut cells have been identified, the next task is to determine which cells of the background discretization belong to the interior of the physical domain and which belong to the exterior. This step is one of the most important components of the SBM workflow, because it defines the active computational domain on which the problem is effectively solved.

The classification is not limited to the cut cells themselves, but rather aims at constructing a global occupancy description of the background domain. To this end, the procedure combines two levels of information: a local pointwise classification with respect to the immersed surface, and a cell-wise classification based on the sampling of interior test points.

Pointwise inside/outside test

To determine whether a given point belongs to the interior or the exterior of the physical domain, the true boundary surface is used as the geometric reference. Conceptually, the procedure

consists of locating the closest portion of the boundary surface, computing the oriented normal vector at that location, and evaluating the relative position of the point with respect to the side indicated by the normal.

In geometric terms, the oriented normal locally distinguishes the interior from the exterior side of the surface. The sign of the projection of the vector connecting the surface to the point onto the oriented normal then provides the inside/outside information. This local criterion is implemented efficiently by means of spatial search structures, which allow repeated point-to-surface queries at an acceptable computational cost.

Use of fake Gauss points for cell classification

The classification of a cut cell is not based on a single point evaluation, but on the sampling of multiple points distributed inside the cell. These points are often referred to as fake Gauss points, since they play a role analogous to quadrature points, although their present purpose is geometric classification rather than numerical integration.

For each cut knot span, a set of internal sampling points is generated. Each point is classified as inside or outside the physical domain according to the pointwise criterion described above, and the number of interior points is then evaluated. Based on this information, a threshold criterion governed by a parameter, typically denoted by λ , is used to decide whether the cell should remain active or be considered external. In this way, the surrogate domain is not determined solely by geometric intersection, but by a quantitative measure of how much of the cell belongs to the physical domain.

Propagation of the classification to neighboring cells

In addition to the cells directly intersected by the boundary, the method also considers neighboring cells in order to ensure a coherent description of the active domain around the interface. After the cut cells have been marked, adjacent cells are examined and classified, typically by evaluating the position of their center or a small set of internal test points with respect to the physical boundary.

This propagation step serves several purposes. It completes the occupancy information in the vicinity of the interface, avoids isolated inconsistencies in the active domain, and improves the topological regularity of the final computational region. Depending on the configuration, a subsequent topological cleanup may also be performed to remove disconnected numerical islands or other local artifacts generated during the classification stage.

Construction of the surrogate boundary

Once the active and inactive cells have been identified, the surrogate boundary can be constructed. This boundary does not coincide with the true physical boundary. Instead, it is defined as the boundary of the active computational domain induced by the background discretization.

In three dimensions, the surrogate boundary is given by the union of the faces of those knot spans that separate an active interior cell from an inactive exterior one. Geometrically, it is therefore a staircase-like surface aligned with the background grid. Numerically, however, it plays a central role, since it is precisely on this surrogate boundary that the SBM boundary conditions are imposed.

This reflects the main idea of the Shifted Boundary Method: the imposition of the boundary conditions is shifted from the true, nonconforming physical boundary to a surrogate boundary that is computationally convenient, while the formulation is corrected in a consistent manner to account for the discrepancy between the two.

Construction of geometric entities for integration

After the surrogate boundary has been extracted, it is transformed into a geometric representation compatible with the IGA data structures used in the Kratos workflow. This allows the background volume, the surrogate boundary, and the corresponding integration entities to be handled within a unified geometric framework.

The background NURBS volume remains the reference geometry for the volumetric elements, while the surrogate boundary is organized in such a way that quadrature-point geometries can be generated for the assembly of the SBM boundary conditions. As a consequence, both the volume elements and the boundary conditions are treated as geometric entities on which shape functions, derivatives, and integration operators can be evaluated in a consistent manner.

Projection from the surrogate boundary to the true boundary

A crucial stage of the workflow is the construction of the geometric correspondence between the surrogate boundary and the true physical boundary. For each integration entity associated with the surrogate boundary, and in particular for the quadrature points used in the boundary conditions, the corresponding location on the true boundary is identified.

Conceptually, the procedure consists of taking a point on the surrogate boundary, locating the closest portion of the true surface, and storing the corresponding closest-point projection data, together with the local geometric information associated with the boundary. This information is then used to construct the SBM correction terms that permit boundary conditions to be imposed on the surrogate boundary while retaining consistency with the physical boundary.

It is precisely through this projection and the associated distance information that the “shift” of the method is introduced: the numerical boundary and the physical boundary do not coincide, but the formulation compensates for their separation by incorporating the geometric offset into the discrete boundary operators.

Creation of volume elements and SBM boundary conditions

Once the geometric workflow has been completed, the numerical entities required for the discrete problem are generated. The volume elements are created on the active portion of the IGA background grid, that is, on the cells classified as belonging to the computational domain.

The SBM boundary conditions are instead created on the surrogate boundary and enriched with the projection information linking them to the true boundary.

As a result, the governing equations are solved over the background IGA domain, whereas the physical boundary conditions are not imposed directly on the immersed surface, but rather on the surrogate boundary. The correction introduced by the SBM formulation ensures that the resulting problem remains consistent with the actual boundary conditions posed on the true physical geometry.

Summary of the workflow

The overall three-dimensional SBM-IGA workflow implemented in Kratos Multiphysics may therefore be summarized as follows. The true geometry is used as an immersed geometric reference, while the computational domain is constructed over a regular IGA background discretization. The cells intersected by the boundary are identified from the imported surface representation, and the active domain is determined through inside/outside tests and cell-wise sampling based on fake Gauss points. From this occupancy information, a surrogate boundary aligned with the background grid is extracted. Finally, the boundary conditions are imposed on this surrogate boundary by exploiting a projection onto the true boundary, thereby preserving physical consistency and numerical accuracy.

This approach combines the geometric richness and smoothness properties of IGA with the flexibility of immersed methods, while eliminating the need for a fully body-fitted three-dimensional mesh in the presence of geometrically complex domains.

4.1 Turek CFD benchmark

Beyond the elliptic and Stokes-type problems addressed in the previous sections, the developed framework was further extended to transient incompressible Navier–Stokes simulations within the open-source platform **Kratos Multiphysics**. This additional development was not part of the three core journal publications, but it plays an important role in demonstrating the broader applicability of the SBM-IGA technology to time-dependent CFD workflows. In particular, it shows that the cut-cell-free immersed philosophy can be integrated into a transient solver environment while retaining the geometric flexibility already observed in the static and steady

$$\rho \left(\frac{\partial \mathbf{u}}{\partial t} + \mathbf{u} \cdot \nabla \mathbf{u} \right) - \nabla \cdot \boldsymbol{\sigma}(\mathbf{u}, p) = \mathbf{f} \quad \text{in } \Omega,$$

$$\nabla \cdot \mathbf{u} = 0 \quad \text{in } \Omega.$$

problems.

$$\boldsymbol{\sigma}(\mathbf{u}, p) = 2\mu \boldsymbol{\varepsilon}(\mathbf{u}) - p \mathbf{I},$$

$$\boldsymbol{\varepsilon}(\mathbf{u}) = \frac{1}{2} (\nabla \mathbf{u} + \nabla \mathbf{u}^T).$$

The resulting workflow is based on a transient Navier–Stokes formulation with SBM-IGA discretization in two and three dimensions. The implementation was carried out in Kratos Multiphysics and validated on the classical **Turek CFD benchmark**, which is widely used as a reference problem for incompressible flow past a bluff obstacle. In this test, the objective is not only to recover a qualitatively correct wake pattern, but also to reproduce the expected integral quantities, in particular the time evolution of drag and lift.




Figure 18: Turek CFD benchmark: problem geometry.




Figure 19: Turek CFD benchmark: The SBM approach: true and surrogate boundaries.

The figure contains a field-level comparison of the immersed Shifted Boundary Method–Isogeometric Analysis (SBM-IGA) solution against a body-fitted finite element reference. Specifically, it displays the velocity and pressure contours. The purpose of this comparison is to show the qualitative agreement between the SBM-IGA discretization and the reference, confirming that the immersed method can reproduce the characteristic flow features of the benchmark without needing a body-fitted mesh.

A first comparison was performed at the field level by contrasting the immersed SBM-IGA solution against a body-fitted finite element reference. The velocity and pressure contours show a very good qualitative agreement, confirming that the immersed discretization is able to reproduce the characteristic flow features of the benchmark without requiring a body-fitted mesh. This is particularly relevant in the context of design-through-analysis workflows, where avoiding repeated remeshing is one of the main motivations for embedded approaches.




Figure 20: Turek CFD benchmark: field-level comparison between the immersed SBM-IGA solution and a body-fitted finite element reference. The velocity and pressure contours show very good qualitative agreement, confirming that the embedded formulation reproduces the characteristic flow features of the benchmark without requiring a body-fitted mesh.

A second and more quantitative comparison is obtained from the drag and lift histories. The results show that the SBM-IGA solutions follow the expected reference trends, and that increasing the spatial resolution improves the agreement with the benchmark curves. These tests therefore provide a first validation of the transient SBM-IGA workflow in an unsteady CFD setting and support its use as a basis for future developments toward more challenging fluid-structure and moving-boundary problems.




Figure 21: Turek CFD benchmark: drag and lift histories for $t \in [9, 9.6]$. The left panel shows the drag force and the right panel the lift force. The red curve corresponds to the Nek5000 reference solution, the black curve to the Turek and Hron reference data, the orange curve to the coarsest SBM-IGA discretization (44,103 DOFs), the blue curve to an intermediate SBM-IGA discretization, and the green curve to the finest SBM-IGA discretization (147,303 DOFs). The results show that increasing the spatial resolution improves the agreement with the benchmark curves.

Overall, the Turek benchmark confirms that the developed Kratos-based workflow extends the SBM-IGA framework beyond model problems and steady Stokes flows, and that the method can be deployed in practical transient CFD simulations while preserving the advantages of immersed high-order discretizations.

4.2 Three-dimensional potentiality

A further objective of this work was to assess the practical potential of the method in three-dimensional configurations. Although the most advanced developments of the thesis, in particular the Gap-SBM multipatch framework, are still limited to two dimensions, the underlying SBM-IGA technology already shows a natural and robust extension path to three-dimensional

problems. This is one of the key practical advantages of the method: once the surrogate-boundary machinery and Taylor-based transfer operators are available, the extension to three-dimensional embedded geometries remains conceptually straightforward.

To illustrate this potential, several three-dimensional flow examples were produced within the same Kratos-based framework. A first example considers flow around a sphere, comparing a classical body-fitted finite element discretization with the corresponding SBM-IGA embedded treatment. The figure shows that the embedded approach is able to reproduce the same basic flow configuration while avoiding the need for a boundary-fitted mesh around the obstacle.




Figure 22: Three-dimensional flow around a sphere: comparison between a body-fitted finite element discretization and the corresponding SBM-IGA embedded treatment. The example illustrates the straightforward extension of the method to three-dimensional configurations.

A second example involves flow around a three-dimensional Stanford bunny. This example is particularly useful because it combines a geometrically complex obstacle with a fully three-dimensional computational setting. The resulting velocity and pressure fields demonstrate that the framework can already handle nontrivial embedded geometries in 3D and therefore provides a promising basis for more realistic industrial applications.




Figure 23: Three-dimensional embedded flow around a Stanford bunny. The result demonstrates the ability of the SBM-IGA framework to handle geometrically complex obstacles in fully three-dimensional settings.

Finally, a third example considers flow around a car-like geometry, again comparing a body-fitted reference with the corresponding SBM-IGA embedded configuration. Although these examples are primarily demonstrative rather than benchmarked in the same systematic manner

as the Turek case, they are important because they show the practical geometric flexibility of the method in three dimensions and its compatibility with complex obstacle shapes.




Figure 24: Three-dimensional flow around a car-like geometry: comparison between body-fitted FEM and SBM-IGA. This example highlights the practical geometric flexibility of the embedded formulation for realistic three-dimensional shapes.

Taken together, these examples support one of the broader conclusions of the thesis: the SBM can be integrated into IGA in a way that is not only cut-cell-free, high-order accurate, and well-conditioned, but also naturally extensible to three-dimensional embedded CFD applications. The present results therefore represent an important intermediate step toward more advanced developments such as moving boundaries, fluid-structure interaction, and fully three-dimensional multipatch coupling.

CONCLUSIONS

This report has presented substantial progress in the development of high-order unfitted isogeometric methodologies based on the Shifted Boundary Method (SBM), addressing key challenges in CFD and related diffusion problems. The work has demonstrated that SBM integrated with Isogeometric Analysis (IGA) provides an accurate and robust alternative to boundary-fitted and trimming-based approaches, enabling simulations on complex geometries through surrogate boundaries and high-order Taylor reconstructions. This strategy avoids cut-cell integration, mitigates the small cut-cell problem, and yields well-conditioned linear systems that are suitable for efficient iterative solvers, while preserving optimal convergence in smooth regimes.

For elliptic problems, the report has confirmed the effectiveness of high-order SBM-IGA formulations, including penalty-free Nitsche enforcement and optimal surrogate boundary selection, achieving stable and optimally convergent solutions without penalty calibration. For incompressible viscous flows, the SBM-IGA framework has been extended to the Stokes equations through Variational Multi-Scale (VMS) stabilization, enabling equal-order velocity–pressure discretizations while maintaining numerical stability and accuracy in embedded settings. The methodology has further been applied to non-Newtonian and viscoplastic models, including Bingham-type fluids, highlighting how spline smoothness interacts with yield surfaces and non-smooth solution features. These results clarify the benefits and limitations of highly regular spline spaces in viscoplastic flow simulations and motivate the use of stabilization and appropriate regularization strategies.

A major recent contribution is the extension of the SBM paradigm to multipatch and nonconforming isogeometric configurations through the Gap–Shifted Boundary Method (Gap–SBM). This development enables high-order, penalty-free coupling between independently parameterized patches with arbitrary refinement, degree, and parametrization, without introducing additional degrees of freedom. The resulting multipatch formulation preserves favorable conditioning and provides a practical mechanism for embedded local refinement and complex geometric assembly, which are central requirements in industrial CAD-based workflows.

An additional relevant outcome of this work is the development of a consistent three-dimensional SBM-IGA workflow within the Kratos Multiphysics framework. This implementation provides the practical infrastructure required to transfer the proposed formulations from the methodological level to realistic computational settings. The workflow separates the physical boundary description from the analysis-suitable background discretization and includes the import of the immersed skin, the construction of the NURBS background grid, the identification of cut knot spans, the inside/outside classification of the background cells, the extraction of the surrogate boundary, and the generation of the corresponding integration entities, support conditions, and active computational elements. This pipeline makes it possible to handle three-dimensional immersed geometries in a manner that remains fully consistent with the SBM philosophy, while preserving the flexibility of IGA discretizations and avoiding the need for body-fitted volumetric meshing. In this sense, the 3D workflow represents a key step toward the reliable application of SBM-IGA to realistic engineering problems and toward its integration into industrial simulation environments.

Future work will focus on extending these developments to the full incompressible Navier–Stokes equations, including nonlinear convection-dominated regimes, and to transient and moving-boundary problems. Another priority is the systematic extension to three-dimensional geometries, particularly for multipatch assemblies and embedded interfaces, including scalable solver and preconditioning strategies. Further developments will also target coupled multiphysics scenarios relevant to industrial CFD, such as fluid–structure interaction and strongly nonlinear rheologies, while maintaining robustness with respect to complex CAD inputs.

Overall, this work establishes a unified and scalable foundation for high-order unfitted isogeometric analysis in CFD, bridging advanced numerical formulations with practical industrial simulation needs and supporting the reduction of preprocessing effort while maintaining high-fidelity predictive capabilities.

REFERENCES

- [1] Thomas J.R. Hughes, John A. Cottrell, Yuri Bazilevs, Isogeometric analysis: CAD, finite elements, NURBS, exact geometry and mesh refinement, *Comput. Methods Appl. Mech. Engrg.* 194 (39) (2005) 4135–4195.
- [2] John A. Cottrell, Alessandro Reali, Yuri Bazilevs, Thomas J.R. Hughes, Isogeometric analysis of structural vibrations, *Comput. Methods Appl. Mech. Engrg.* 195 (41) (2006) 5257–5296.
- [3] Yuri Bazilevs, Lourenco Beirao da Veiga, John A. Cottrell, Thomas J.R. Hughes, Giancarlo Sangalli, Isogeometric analysis: Approximation, stability and error estimates for h-refined meshes, *Math. Models Methods Appl. Sci.* 16 (07) (2006) 1031–1090.
- [4] Yuri Bazilevs, Victor M. Calo, Yongjie Zhang, Thomas J.R. Hughes, Isogeometric fluid–structure interaction analysis with applications to arterial blood flow, *Comput. Mech.* 38 (4) (2006) 310–322.
- [5] Yongjie Zhang, Yuri Bazilevs, Samrat Goswami, Chandrajit L. Bajaj, Thomas J.R. Hughes, Patient-specific vascular NURBS modeling for isogeometric analysis of blood flow, *Comput. Methods Appl. Mech. Engrg.* 196 (29) (2007) 2943–2959.
- [6] John A. Cottrell, Thomas J.R. Hughes, Alessandro Reali, Studies of refinement and continuity in isogeometric structural analysis, *Comput. Methods Appl. Mech. Engrg.* 196 (41) (2007) 4160–4183.
- [7] Josef Kiendl, Kai-Uwe Bletzinger, Johannes Linhard, Roland Wüchner, Isogeometric shell analysis with Kirchhoff–Love elements, *Comput. Methods Appl. Mech. Engrg.* 198 (49) (2009) 3902–3914.
- [8] David J. Benson, Yuri Bazilevs, Ming-Chen Hsu, Thomas J.R. Hughes, Isogeometric shell analysis: The Reissner–Mindlin shell, *Comput. Methods Appl. Mech. Engrg.* 199 (5) (2010) 276–289, *Computational Geometry and Analysis*.
- [9] Kenji Takizawa, Yuri Bazilevs, Tayfun E. Tezduyar, Ming-Chen Hsu, Takuya Terahara, Computational cardiovascular medicine with isogeometric analysis, *J. Adv. Eng. Comput.* 6 (3) (2022) 167–199.
- [10] Massimo Carraturo, Carlotta Giannelli, Alessandro Reali, Rafael Vázquez, Suitably graded THB-spline refinement and coarsening: Towards an adaptive isogeometric analysis of additive manufacturing processes, *Comput. Methods Appl. Mech. Engrg.* 348 (2019) 660–679.
- [11] Carlotta Giannelli, Bert Jüttler, Hendrik Speleers, THB-splines: The truncated basis for hierarchical splines, *Comput. Aided Geom. Design* 29 (7) (2012) 485–498, *Geometric Modeling and Processing 2012*.
- [12] Dominik Schillinger, Luca Dedè, Michael A. Scott, John A. Evans, Michael J. Borden, Ernst Rank, Thomas J.R. Hughes, An isogeometric design-through-analysis methodology based on adaptive hierarchical refinement of NURBS, immersed boundary methods, and T-spline CAD surfaces, *Comput. Methods Appl. Mech. Engrg.* 249–252 (2012) 116–150, *Higher Order Finite Element and Isogeometric Methods*.

- [13] Martin Ruess, Dominik Schillinger, Yuri Bazilevs, Vasco Varduhn, Ernst Rank, Weakly enforced essential boundary conditions for NURBS-embedded and trimmed NURBS geometries on the basis of the finite cell method, *Internat. J. Numer. Methods Engrg.* 95 (10) (2013) 811–846.
- [14] Jamshid Parvizian, Alexander Düster, Ernst Rank, Finite cell method: h-and p-extension for embedded domain problems in solid mechanics, *Comput. Mech.* 41 (1) (2007) 121–133.
- [15] Ernst Rank, Martin Ruess, Stefan Kollmannsberger, Dominik Schillinger, Alexander Düster, Geometric modeling, isogeometric analysis and the finite cell method, *Comput. Methods Appl. Mech. Engrg.* 249–252 (2012) 104–115.
- [16] Michael Breitenberger, Andreas Apostolatos, Philipp Bucher, Roland Wüchner, Kai-Uwe Bletzinger, Analysis in computer aided design: Nonlinear isogeometric B-Rep analysis of shell structures, *Comput. Methods Appl. Mech. Engrg.* 284 (2015) 401–457, Isogeometric Analysis Special Issue.
- [17] Tobias Teschemacher, Anna M. Bauer, Thomas Oberbichler, Micheal Breitenberger, Riccardo Rossi, Roland Wüchner, Kai-Uwe Bletzinger, Realization of CAD-integrated shell simulation based on isogeometric B-Rep analysis, *Adv. Model. Simul. Eng. Sci.* 5 (2018) 1–54.
- [18] Tobias Teschemacher, Anna M. Bauer, Ricky Aristio, Manuel Meßmer, Roland Wüchner, Kai-Uwe Bletzinger, Concepts of data collection for the CAD-integrated isogeometric analysis, *Eng. Comput.* 38 (6) (2022) 5675–5693.
- [19] Manuel Meßmer, Tobias Teschemacher, Lukas F. Leidinger, Roland Wüchner, Kai-Uwe Bletzinger, Efficient CAD-integrated isogeometric analysis of trimmed solids, *Comput. Methods Appl. Mech. Engrg.* 400 (2022) 115584.
- [20] Manuel Meßmer, Stefan Kollmannsberger, Roland Wüchner, Kai-Uwe Bletzinger, Robust numerical integration of embedded solids described in boundary representation, *Comput. Methods Appl. Mech. Engrg.* 419 (2024) 116670.
- [21] Erik Burman, Ghost penalty, *C. R. Math.* 348 (21–22) (2010) 1217–1220.
- [22] Santiago Badia, Eric Neiva, Francesc Verdugo, Linking ghost penalty and aggregated unfitted methods, *Comput. Methods Appl. Mech. Engrg.* 388 (2022) 114232.
- [23] Alex Main, Guglielmo Scovazzi, The shifted boundary method for embedded domain computations. Part I: Poisson and Stokes problems, *J. Comput. Phys.* 372 (2018) 972–995.
- [24] Alex Main, Guglielmo Scovazzi, The shifted boundary method for embedded domain computations. Part II: Linear advection–diffusion and incompressible Navier–Stokes equations, *J. Comput. Phys.* 372 (2018) 996–1026.
- [25] Nabil M. Atallah, Guglielmo Scovazzi, Nonlinear elasticity with the shifted boundary method, *Comput. Methods Appl. Mech. Engrg.* 426 (2024) 116988.

- [26] Efthymios N. Karatzas, Giovanni Stabile, Leo Nouveau, Guglielmo Scovazzi, Gianluigi Rozza, A reduced-order shifted boundary method for parametrized incompressible Navier–Stokes equations, *Comput. Methods Appl. Mech. Engrg.* 370 (2020) 113273.
- [27] Nabil M. Atallah, Claudio Canuto, Guglielmo Scovazzi, The second-generation shifted boundary method and its numerical analysis, *Comput. Methods Appl. Mech. Engrg.* 372 (2020) 113341.
- [28] Nabil M. Atallah, Claudio Canuto, Guglielmo Scovazzi, Analysis of the shifted boundary method for the Poisson problem in domains with corners, *Math. Comp.* 90 (331) (2021) 2041–2069.
- [29] Haydel Collins, Alexei Lozinski, Guglielmo Scovazzi, A penalty-free shifted boundary method of arbitrary order, *Comput. Methods Appl. Mech. Engrg.* 417 (2023) 116301, A Special Issue in Honor of the Lifetime Achievements of T. J. R. Hughes.
- [30] Cheng-Hau Yang, Kumar Saurabh, Guglielmo Scovazzi, Claudio Canuto, Adarsh Krishnamurthy, Baskar Ganapathysubramanian, Optimal surrogate boundary selection and scalability studies for the shifted boundary method on octree meshes, *Comput. Methods Appl. Mech. Engrg.* 419 (2024) 116686.
- [31] Antonelli, N., Aristio, R., Gorgi, A., Zorrilla, R., Rossi, R., Scovazzi, G., Wüchner, R. (2024). The Shifted Boundary Method in Isogeometric Analysis. *Computer Methods in Applied Mechanics and Engineering*, 430, 117228.
- [32] Codina, R., Badia, S., Baiges, J., & Principe, J. (2018). Variational multiscale methods in computational fluid dynamics. *Encyclopedia of computational mechanics*, 1-28.
- [33] Codina, R. (2000). Stabilization of incompressibility and convection through orthogonal sub-scales in finite element methods. *Computer methods in applied mechanics and engineering*, 190(13-14), 1579-1599.
- [34] Mitsoulis E, Zisis T. Flow of Bingham plastics in a lid-driven square cavity. *Journal of Non-Newtonian Fluid Mechanics* 2001; 101(1–3):173 – 180.
- [35] Cotela Dalmau, Jordi, Riccardo Rossi, and Antonia Larese De Tetto. "Simulation of two-and three-dimensional viscoplastic flows using adaptive mesh refinement." *International journal for numerical methods in engineering* (2017).
- [36] Benjamin Marussig and Thomas JR Hughes. A review of trimming in isogeometric analysis: challenges, data exchange and simulation aspects. *Archives of computational methods in engineering*, 25(4):1059–1127, 2018
- [37] Ulrich Langer, Angelos Mantzaflaris, Stephen E Moore, and Ioannis Touloupoulos. Multipatch discontinuous galerkin isogeometric analysis. In *Isogeometric Analysis and Applications 2014*, pages 1–32. Springer, 2015.
- [38] Chiu Ling Chan, Cosmin Anitescu, and Timon Rabczuk. Isogeometric analysis with strong multipatch c1-coupling. *Computer Aided Geometric Design*, 62:294–310, 2018
- [39] Mika Juntunen and Rolf Stenberg. Nitsche's method for general boundary conditions. *Mathematics of computation*, 78(267):1353–1374, 2009.

- [40] Hirschler, Robin Bouclier, A Duval, T Elguedj, and Joseph Morlier. The embedded isogeometric Kirchhoff–Love shell: from design to shape optimization of non-conforming stiffened multipatch structures. *Computer Methods in Applied Mechanics and Engineering*, 349:774–797, 2019.
- [41] Haydel Collins, Kangan Li, Alexei Lozinski, and Guglielmo Scovazzi. Gap-sbm: A new conceptualization of the shifted boundary method with optimal convergence for the neumann and dirichlet problems. *arXiv preprint arXiv:2508.09613*, 2025.
- [42] Rubén Zorrilla, Riccardo Rossi, Guglielmo Scovazzi, Claudio Canuto, and Antonio Rodríguez-Ferran. A shifted boundary method based on extension operators. *Computer Methods in Applied Mechanics and Engineering*, 421:116782, 2024.