Gecko

Design for IGA-type discretization workflows

Gecko Technical Report 2

DC6 – Lucas Venta Viñuela

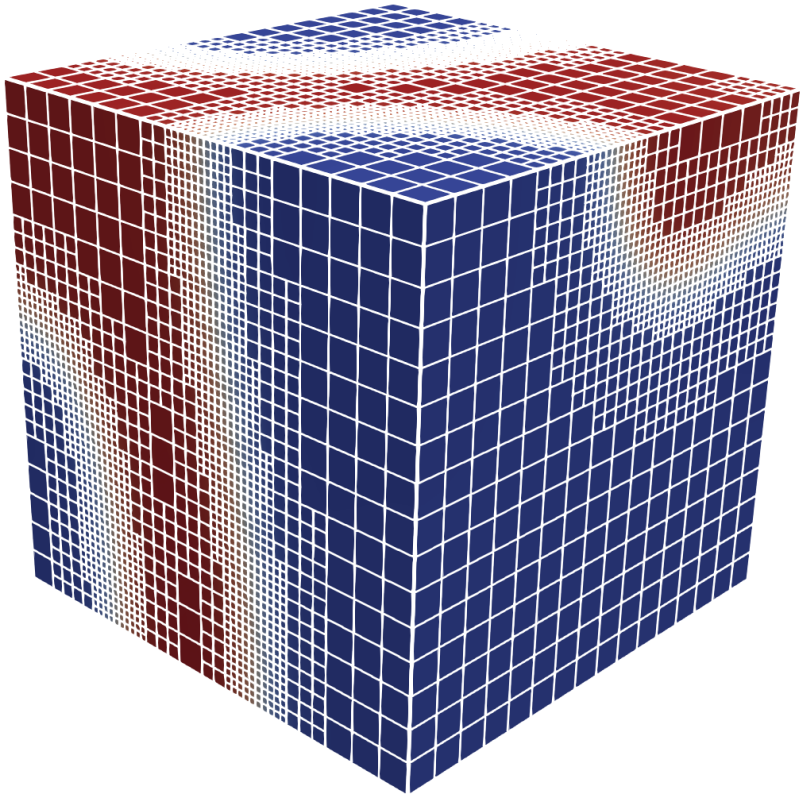

This project has received funding from the European Union's Horizon Europe research and innovation programme under grant agreement No 101073106
Call: HORIZON-MSCA-2021-DN-01

Funded by the European Union

Executive summary

This report presents the progress of my doctoral research within the Gecko Doctoral Network, focusing on the development of adaptive methods for phase-field modeling of interfacial phenomena. The work specifically addresses the numerical simulation of phase separation processes governed by the Cahn–Hilliard equation, with an emphasis on achieving high accuracy and computational efficiency in complex and evolving domains.

The research is conducted within the framework of Isogeometric Analysis (IGA), which enables the use of smooth, spline-based basis functions for the discretization of high-order partial differential equations. This feature is particularly advantageous for higher-order phase-field formulations, where fourth-order spatial derivatives arise naturally. By leveraging the inherent regularity of the shape functions, the proposed approach avoids the need for auxiliary field variables and enables accurate solutions with reduced numerical complexity and improved computational times.

A key component of the work is the development of adaptive meshing strategies based on hierarchical spline constructions, such as truncated hierarchical B-splines (THB-splines). These approach allows for local mesh refinement and coarsening driven by the evolution of the solution, concentrating computational effort in regions characterized by steep gradients, diffuse interfaces, or evolving microstructures. This results in significant gains in efficiency while preserving high-order accuracy.

The methodology is first assessed via a convergence study employing a manufactured solution. Subsequently, numerical examples with randomly generated initial conditions are presented to demonstrate its capability to accurately capture phase evolution and interface dynamics. Particular attention is devoted to convergence behavior, the robustness of the adaptive scheme, and the trade-off between accuracy and computational cost.

Overall, this research contributes to advancing the state of the art in adaptive phase-field simulations by combining high-order modelling and efficient refinement strategies. The proposed framework is especially well suited for phase separation and coarsening processes governed by the Cahn–Hilliard equation, while also offering strong potential for extension to higher-order phase-field formulations in fracture mechanics.

Future work will be devoted to the extension of the proposed framework to three-dimensional simulations, its application to phase-field fracture models, and the development of adaptive immersed methods. These efforts are expected to further enhance the flexibility and applicability of the approach, while strengthening the integration between computational analysis and geometric design tools.

Table of contents

List of figures 4
List of abbreviations 5
Introduction 6
1. Phase-field modelling of the Cahn—Hilliard equation 8
2. Local adaptive meshing in isogeometric analysis 9
2.1. Element marking 10
2.2. Transfer of the field variables 11
3. Convergence analysis based on a manufactured solution 11
4. Phase separation numerical examples 14
4.1. Nucleation-driven phase separation in 2D 15
5. Conclusions and outlook 18
6. References 19

List of figures

Figure 1.-Schematic illustration of the adaptive isogeometric loop. 9
Figure 2.- Element marking based on a solution field using a threshold. (a) Initial mesh showing the solution field, with yellow contours indicating the threshold used for marking elements. (b) Elements marked for refinement. (c) Elements marked for coarsening. 10
Figure 3.- Manufactured solution used for the convergence study, featuring a smooth, radially symmetric interface centered at $xc, yc = 0.5,0.5$, shown at the initial time $t = 0\ s$.....13
Figure 4.- Errors measured in the $L2$-, $H1$-, and $H2$-norms, shown in (a)–(c) for quadratic ($p = 2$), (d)–(f) for cubic ($p = 3$), and (g)–(i) for quartic ($p = 4$) basis functions, plotted against the number of degrees of freedom (DOFs) for both tensor product (TP) and truncated hierarchical B-spline (THB) meshes. Here, u denotes the manufactured (exact) solution and uh the numerical solution. 14
Figure 5.- Nucleation in 2D. Evolution of the concentration from an initially perturbed state $ux, 0 = u + \delta ux$, with $u = 0.4$ and δux a small random perturbation uniformly distributed in $-0.005,0.005$. Top row (a-c): tensor product (TP) mesh; middle row (d-f): adaptivity using $L2$-projection (THB L2); bottom row (g-i): adaptivity using the quasi-interpolant based on local interpolation (THB QI). The mesh consists of 642 quadratic elements in the TP construction. 16
Figure 6.- Overview of the computational insights during nucleation in 2D, comparing three approaches — tensor product (TP), adaptivity with $L2$-projection (THB L2), and adaptivity with quasi-interpolant (THB QI) — for the evolution of degrees of freedom (DOFs) (a), total energy $\mathcal{G}$(b), and mass difference $\Delta m = \Omega u\, d\Omega - \Omega u\mathrm{TP}\, d\Omega$ (c), where $u\mathrm{TP}$ is the solution computed with the TP mesh. Computational times are reported for assembly time (AT, d), solver time (ST, e), and projection time (PT, f). The initial condition is $ux, 0 = u0 + \delta ux$ with $u0 = 0.4$ and δux a small random perturbation uniformly distributed in $-0.005,0.005$. The mesh consists of 642 quadratic elements in the TP construction.....17
Figure 7.- Total computational times for 2D nucleation using three approaches: tensor product (TP), adaptivity with $L2$-projection (THB L2), and adaptivity with the quasi-interpolant (THB QI). Two refinement strategies are compared: (a) simulations without a maximum number of refinement iterations, and (b) simulations where refinement iterations are triggered by significant changes in the total energy.....17

List of abbreviations

CAE	*Computer-Aided Engineering*
FCM	*Finite Cell Method*
FEM	*Finite Element Method*
IGA	*IsoGeometric Analysis*

Introduction

Phase separation phenomena are ubiquitous in nature, occurring in a wide range of systems, including metallic alloys, polymer blends, and biological membranes. These processes involve the spontaneous separation of a homogeneous mixture into distinct phases, driven by the minimization of the free energy of the system. The Cahn--Hilliard equation [1,2] provides a canonical mathematical framework for phase separation, describing the evolution of a conserved order parameter, typically the concentration field, through a fourth-order partial differential equation. Its governing dynamics couple diffusion with interfacial energy contributions, naturally capturing the evolution of the interfaces. The fourth-order nature of the Cahn--Hilliard equation requires numerical discretizations that are at least C^1-continuous in the weak form.

Classical finite element methods are typically only C^0-continuous and cannot directly satisfy this requirement and therefore often employ mixed formulations [3], introducing an auxiliary variable to solve a coupled system of two second-order PDEs. While this enables the use of standard finite elements, it increases the number of unknowns and the complexity of the system. Instead, finite difference schemes implement central difference approximations for the Laplacian and bi-Laplacian operators, providing a simple and efficient solution on uniform grids. However, since these schemes rely on structured, evenly spaced meshes, they are limited in handling complex geometries or curved boundaries [4]. Finally, discontinuous Galerkin (DG) methods [5] allow higher-order discontinuous approximations to discretize the fourth-order problem but introduce additional degrees of freedom and require stabilization strategies.

An alternative approach that addresses both the continuity requirements of fourth-order problems and the challenges of using complex geometries is isogeometric analysis (IGA) [6], which has enabled the widespread use of high-order spline spaces for the numerical discretization of partial differential equations. The core idea of IGA is to use the same basis functions used in computer-aided design (CAD), namely B-splines and non-uniform rational B-splines (NURBS), to represent both geometry and solution spaces. Beyond preserving the geometry representation from the CAD, these bases provide higher-order global continuity, which follows directly from their construction. This property is particularly relevant for the numerical treatment of higher-order models -- such as those arising in thin-plate and shell theories [7,8], phase-field formulations [9,10,11,12], and flow and transport simulations [13,14] -- all involving higher-order differential operators. In contrast to classical finite element spaces, spline-based spaces inherently provide the necessary smoothness and can be constructed with arbitrary high-order, enabling a direct and efficient discretization of higher-order PDEs while preserving the CAD representation of the geometry.

Phase-field modeling has emerged as a powerful framework for simulating interface problems. In these models, sharp interfaces are replaced by thin transition regions where the order parameter varies smoothly between phases. The width of these regions is governed by a length-scale parameter: as this parameter decreases, the diffuse interface approaches a sharp-interface description, but the resulting steep gradients necessitate high spatial resolution. A straightforward strategy is to employ uniform global h-refinement until the length scale is adequately resolved, but this quickly becomes computationally prohibitive. A more efficient alternative is to use adaptive methods, which enable local refinement and coarsening of the mesh in response to the evolving phase-field [15]. We refer to [16] and references therein for a more general overview of the mathematical foundations of adaptive isogeometric methods.

While B-spline bases are fundamental to isogeometric analysis, their standard tensor-product construction [17] presents a significant limitation for adaptive refinement. A key drawback is that refining a single element using tensor-product constructions can propagate along an entire direction, needlessly adding many degrees of freedom in areas outside the region of interest. To address this issue, there

are several spline constructions that allow for local refinement in the literature. Popular constructions include hierarchical B-splines (HB-splines) [18], truncated hierarchical B-splines (THB-splines) [19,20], T-splines and their variations [21,22,23], polynomial splines over hierarchical T-meshes (PHT-splines) [24], locally refined (LR) splines [25], and splines over unstructured meshes (U-splines) [16]. Among them, HB-splines are arguably the easiest to define and implement within the IGA framework, thanks to their multilevel structure [26]. THB-splines preserve all the desirable properties of HB-splines while offering smaller support and forming a partition of unity due to truncation, making them particularly well-suited for numerical analysis. Importantly, THB-splines naturally extend across dimensions, allowing straightforward adaptive refinement and coarsening in both 2D and 3D simulations.

Although adaptive phase-field simulations have been investigated previously [15,27,28], these studies were typically limited to two-dimensional problems, and efficiency assessments were often restricted to reductions in the number of degrees of freedom. In contrast, our goal is to develop a framework that is not only adaptive but also computationally faster in both 2D and 3D. While THB-splines introduce some overhead compared to uniform tensor-product bases, this work assesses whether their use, combined with adaptive refinement and quasi-interpolation [29,30] for the projection of solution fields during refinement/coarsening, can lead to overall speed-ups in realistic phase-field simulations.

In this work, we develop an adaptive isogeometric analysis framework for phase-field modeling of phase separation, capable of handling both two- and three-dimensional problems, and demonstrating clear potential for reducing computational cost. The approach leverages the Cahn--Hilliard equation, whose fourth-order differential operator is naturally discretized using B-splines of degree p, providing C^{p-1} global continuity across the knot spans. Moreover, high-regularity splines present superior approximation properties of such spline functions on a per-degree-of-freedom basis [31]. Our choice for using an adaptive framework with THB-splines is motivated by the need to efficiently resolve and track the interfaces inherent to the Cahn--Hilliard dynamics. The phase-field method naturally generates sharp but smooth transition layers between phases, which must be accurately captured to preserve the physical properties of the model. Uniform refinement is prohibitively expensive in this context, as the interfacial regions typically occupy only a small portion of the computational domain. THB-splines are a natural extension of the HB-spline constructions obtained by locally truncating the basis functions from coarser levels. This truncation recovers the partition of unity, reduces the overlap between basis functions of different refinement levels, and enhances the numerical stability of the resulting discretization. At the same time, THB-splines inherit the desirable properties of standard HB-splines, including linear independence and non-negativity, thereby providing an efficient and robust basis for adaptive isogeometric analysis.

1. Phase-field modelling of the Cahn—Hilliard equation

The Cahn--Hilliard equation [1,2] describes the phase separation process assuming a simplified isotropic and isothermal framework. The thermodynamic state of the mixture is described by a function $u(\mathbf{x}, t)$ of space $\mathbf{x}$ and time t, representing the continuous field variable --the field *order parameter*-- that describes the concentration of the species in the mixture. Under the isothermal hypothesis, the thermodynamic potential for the two-phase immiscible mixture is the Ginzburg-Landau free energy, defined as the functional $\mathcal{G}: H^1(\Omega) \mapsto \mathbb{R}$ (with Ω being an open subset of $\mathbb{R}^d$ and d being the spatial domain dimension):

$$\mathcal{G}[u] = \int_\Omega \left(F(u) + \frac{\lambda}{2} |\nabla u|^2 \right) d\mathbf{x}.$$

The first term of the functional $\mathcal{G}$ defines bulk free energy, a double-well function that allows for the coexistence of two phases, and is a function of the molar concentration per unit volume u. Among the many definitions of the bulk energy, the Flory-Huggins potential [3,4] is often used in phase-separation models. However, more simple polynomial expressions have also been used, some derived from the linearization of the Flory-Huggins potential and others being purely mathematical polynomial expressions. In our model, we define the bulk chemical free energy F as

$$F(u) = \tfrac{1}{4}(1 - u^2)^2,$$

which is a non-convex polynomial of quartic type with a double-well structure characterized by two equilibrium concentrations (i.e., the binodal points) at $u = \pm 1$. The second term of the functional $\mathcal{G}$ accounts for the interfacial free energy of the mixture and is a function of the concentration gradient, representing an energetic penalty for the regularization of the interface, where λ is the interface parameter controlling the thickness of the interface.

The Cahn--Hilliard equation is then obtained by imposing mass conservation, i.e., $\frac{\partial u}{\partial t} + \nabla \cdot \boldsymbol{j} = 0$, assuming a mass flux

$$\boldsymbol{j} = -M(u) \nabla \left(\frac{\delta G}{\delta u} \right),$$

where $M(u)$ is the concentration-dependent mobility, a material-dependent transport relationship that generally depends on the mixture composition and affects the coarsening kinetics [9], determining how fast the concentration field $u(\mathbf{x}, t)$ evolves over time. This ensures that the free energy of the system decreases over time, with $\frac{\delta G}{\delta u}$ representing the variational derivative of the free energy with respect to the variation of the phase-field, δu. Thus, the governing equation reads

$$\frac{\partial u}{\partial t} = \nabla \cdot [M(u) \nabla (F'(u) - \lambda \Delta u)],$$

where $(\cdot)'$ represents the derivative $\partial(\cdot)/\,\partial u$.

The minimization of the free energy of the system drives the dynamics of the solution of the Cahn--Hilliard equation above, with the phase field u approaching the binodal points while creating regions of each pure phase that evolve over time and are separated by smooth and continuous interfaces represented by the phase field, with a thickness of a characteristic length in the order of $\sqrt{\lambda}$.

The strong form of the problem can be written by complementing the Cahn--Hilliard equation with a series of boundary and initial conditions. Let $\Omega \subset \mathbb{R}^{\mathbb{d}}$, be an open set of dimension d that contains the mixture within the domain boundary $\partial\Omega$, assumed to be sufficiently smooth and composed of two

complementary parts on which the Dirichlet and Neumann boundary conditions are specified, $\partial\Omega = \overline{\partial\Omega_D \cup \partial\Omega_N}$, and the unit outward normal to $\partial\Omega$ is denoted as $\boldsymbol{n}$. The initial/boundary value problem over the spatial domain Ω and the time interval $[0,T]$ can then be stated as follows: Find $u: \overline{\Omega} \times [0,T] \to \mathbb{R}$ such that

$$\frac{\partial u}{\partial t} - \nabla \cdot \big(M(u)\nabla(F'(u) - \lambda\Delta u)\big) = 0 \quad \text{in } \Omega \times [0,T]$$

$$u = g \quad \text{on } \partial\Omega_D \times [0,T]$$

$$M(u)\nabla(F'(u) - \lambda\Delta u) \cdot n = 0 \quad \text{on } \partial\Omega_N \times [0,T]$$

$$\nabla u \cdot n = 0 \quad \text{on } \partial\Omega \times [0,T]$$

$$u(x,0) = u_0(x) \quad \text{in } \Omega$$

The Cahn--Hilliard equation is inherently time-dependent; however, this work is concerned primarily with the development of mesh adaptivity techniques to track the evolution of the interfaces efficiently. To maintain the emphasis on this aspect, we use a fixed time-stepping scheme and do not consider adaptivity in the time-step size for the temporal discretization. This choice allows for a clearer assessment of the mesh adaptivity without introducing additional complexity from time step adaptivity effects. For studies addressing the effects of time discretization, the reader is referred to [5,6].

2. Local adaptive meshing in isogeometric analysis

The presence of steep transitions in phase-field models, that need to be resolved by the adopted mesh, constitute a major bottleneck in large-scale simulations. Adaptive meshing schemes offer a remedy for tracking these singular features, allowing to keep the total number of elements reasonable also in large-scale numerical models.

Adaptive algorithms typically follow an iterative loop consisting of four main steps: solve, estimate, mark, and refine/coarsen [20]. In the estimate step, a suitable error estimator provides a quantitative measure of the local approximation error, which is then used in the mark step to identify elements for refinement or coarsening. While in many problems, such as elasticity, this requires carefully designed a-posteriori estimators, in the context of phase-field models the interface location itself can serve as a reliable indicator of local error. Specifically, for the Cahn--Hilliard equation, we adopt a marking strategy where elements are marked based on the magnitude of the numerical solution, which identifies the interface region. This approach has been shown to provide a reliable estimator for adaptive refinement in phase-field simulations [15].

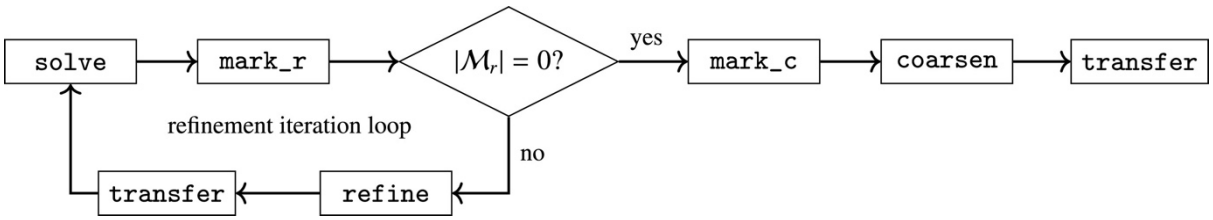

Figure 1.-Schematic illustration of the adaptive isogeometric loop.

In our work, the adaptive isogeometric meshing method follows the iterative procedure illustrated in Figure 1. Each cycle begins with a solve step, followed by mark_r, where elements are marked for refinement and collected in the refinement set $\mathcal{M}_r$ based on the numerical solution u. If $|\mathcal{M}_r| > 0$, the refinement iteration loop --- consisting of refine, transfer, and a return to solve --- is executed until

$|\mathcal{M}_r| = 0$. Once refinement is complete, elements are marked for coarsening using mark_c, coarsened, and the solution is transferred to the updated mesh.

The procedures presented for local adaptive meshing operate in a multidimensional setting, both in the construction of the adopted spline basis and in the meshing process. The adaptive algorithm naturally extends to two- and three-dimensional problems, providing significant flexibility in resolving evolving features. In this section, we present the construction of THB-splines, the element marking strategy, the admissible refinement and coarsening of the mesh, and the transfer of field variables between meshes. Additional details and illustrative examples of these procedures can be found in [7,8,9].

2.1. Element marking

As soon as localized error contributions are known, elements can be marked for refinement or coarsening. This marking is mostly done using the Dörfler marking strategy [37], which involves marking the regions with the largest error contributions until their sum exceeds a certain percentage of the total error. An alternative is to mark the regions with an error higher than a threshold (an absolute threshold based on the maximum error) [19] or based on a relative threshold, taking a fixed percentage of the total number of cells for refinement.

Since the Cahn--Hilliard equation produces a bounded solution $u \in [-1,1]$, the equilibrium phases correspond to regions where the concentration takes the values $u = \pm 1$. Regions with intermediate concentrations indicate the presence of interfaces and therefore require mesh refinement, whereas areas near the equilibrium phases can be coarsened. To accurately capture the interfaces, elements are marked using a threshold-based criterion that distinguishes interface regions from bulk phases. In contrast to residual-based strategies, where local error indicators determine the elements to be refined, our approach performs the marking directly on the solution field u.

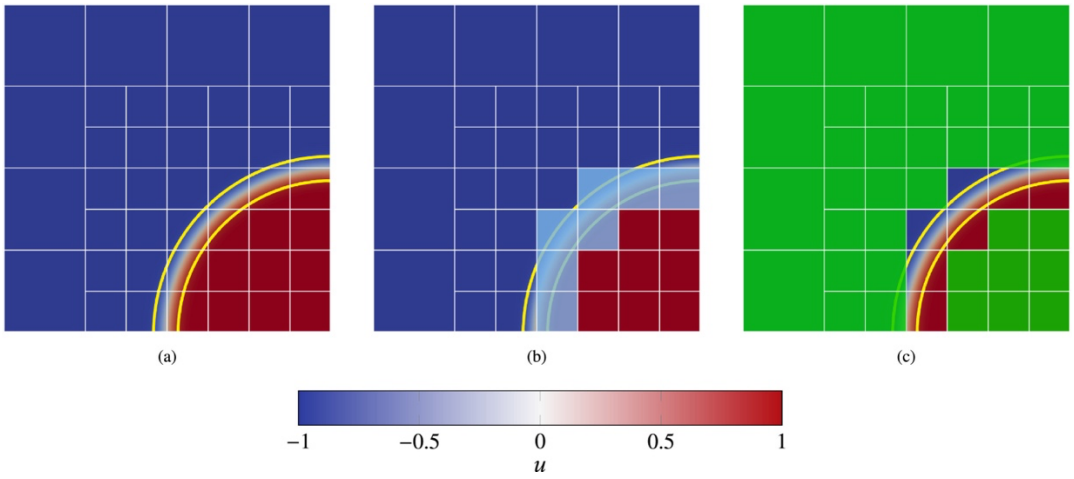

Figure 2.- Element marking based on a solution field using a threshold. (a) Initial mesh showing the solution field, with yellow contours indicating the threshold used for marking elements. (b) Elements marked for refinement. (c) Elements marked for coarsening.

Figure 2 illustrates threshold-based marking for a smooth solution with concentrations $u \in [-1,1]$. For a marking threshold $\Theta = 0.1$, the yellow contours at $u = \pm 0.9$ indicate the concentration values used to determine element refinement or coarsening. Elements with average concentrations within these contours, corresponding to the interface regions, are marked for refinement, while elements outside the contours are marked for coarsening. The figure explicitly highlights the elements to be refined and coarsened. This threshold-based marking accurately resolves interfaces without relying on traditional

error indicators, which is particularly useful for the Cahn--Hilliard equation that lacks an analytical solution.

2.2. Transfer of the field variables

In adaptive isogeometric analysis, the computational mesh is continuously refined or coarsened to accurately capture evolving features of the solution. Such dynamic modifications of the underlying spline space require the solution fields --including concentrations, fluxes, or other internal variables-- to be transferred from the previous mesh to the newly adapted one. This aspect is particularly important in the framework of evolutionary equations such as the Cahn–Hilliard equation, where the solution at a given time step depends directly on the fields computed at the previous step. Ensuring an accurate transfer of these fields is therefore essential: even small inconsistencies can degrade the convergence of the Newton–Raphson method, violate conservation properties, or introduce spurious oscillations that may accumulate over successive steps.

A spline quasi-interpolation scheme constructs an approximation by exploiting the local support of spline basis functions: each coefficient is determined from information in a small neighborhood, and the global approximation is obtained by combining these local contributions. In the hierarchical spline setting, the approximation is organized across multiple refinement levels. Thanks to the coefficient preservation property of truncated hierarchical B-splines, coefficients computed on coarser levels remain valid after refinement, so the hierarchical quasi-interpolant can be assembled by reusing these local contributions across levels.

The coefficients are defined through local functionals, typically based on values (and possibly derivatives or averages) of the target function near the support of each basis function. With a suitable choice, the scheme can exactly reproduce polynomials up to a given degree or act as a projector onto the hierarchical spline space. In the considered approach, coefficients are obtained by solving small local interpolation problems on individual elements, enforcing agreement with the function at selected sampling points (e.g., Gauss points). Since each coefficient depends only on basis functions overlapping a given element, the procedure is fully local and naturally parallelizable.

3. Convergence analysis based on a manufactured solution

The absence of analytical solutions for the Cahn--Hilliard equation complicates quantitative convergence analysis. Furthermore, achieving sufficient accuracy and stability for higher-order PDEs can be computationally intensive, especially for high-dimensional and time-dependent problems, often demanding extremely fine discretizations or the development of advanced numerical schemes. In this context, it is common practice to validate numerical methods through results from highly refined discretizations or manufactured solutions.

In our work, we carry out the convergence study based on a manufactured solution. To choose an analytical reference solution, we choose a function that generally does not satisfy the underlying differential equations. Thus, a residual must be included into the right-hand side of the discrete system. The inherent non-linearity of the equation implies that the optimal convergence rates observed in linear scenarios are primarily illustrative. Nevertheless, given the smoothness of the free energy functional and the absence of discontinuities, we anticipate that these rates will be closely achieved.

Before addressing the convergence of the nonlinear Cahn--Hilliard equation, we recall the classical convergence theory for linear higher-order partial differential equations. A foundational tool in this context is the Aubin--Nitsche argument, which provides a powerful framework for deriving *a priori* error estimates in the lower-order norms. As presented in [38], for a linear elliptic PDE of order 2μ, with $\mu \geq 1$, let r be an integer such that $0 \leq r < \mu$. Let $u \in H^r(\Omega)$, with $r \geq \mu$, be the exact solution of a well-posed scalar elliptic PDE, and let u^h be the approximate solution obtained using B-spline-based isogeometric analysis with basis functions of degree p and minimum regularity $k_{min} \geq \mu$. Then, the following *a priori* error estimate holds:

$$\|u - u^h\|_{H^\sigma(\Omega)} \leq C_\sigma \, h^s \|u\|_{H^r(\Omega)},$$

where $s \coloneqq \min\{\delta - \sigma, 2(\delta - \mu)\}$, with $\delta \coloneqq \min\{p+1, r\}$.

The convergence rate of the error with respect to the global mesh size h depends on the order of the problem m, the regularity r of the exact solution $u \in H^r(\Omega)$,and the degree p of the B-spline basis functions used to represent the approximate solution $u^h \in V^h$. Moreover, when considering the a priori error estimate in lower-order norms, the rate of convergence s depends also on the norm $H^\sigma(\Omega)$ used for the estimation of the error $u - u^h$. For a linear fourth-order PDE $(\mu = 2)$, and an infinitely continuous reference solution

$$|u - u^h|_{H^0 \equiv L^2} \leq C_0 \, h^{\min\{p+1,\, 2p-2\}} |u|_{H^r},$$

$$|u - u^h|_{H^1} \leq C_1 \, h^{\min\{p,\, 2p-2\}} |u|_{H^r},$$

$$|u - u^h|_{H^2} \leq C_2 \, h^{\min\{p-1,\, 2p-2\}} |u|_{H^r}.$$

To enable a meaningful assessment of adaptive refinement strategies, we require a manufactured solution that better reflects the qualitative features of phase-field models, and specifically the Cahn--Hilliard phase separation patterns. The function should saturate at the binodal points ± 1, forming large, nearly uniform bulk phases separated by smooth interfaces represented by the field parameter. This structure facilitates locally refined meshes across the interfaces while allowing coarser meshes elsewhere. Moreover, the function must be time-dependent and sufficiently differentiable in space to enable the analytical computation of the source term in the Cahn--Hilliard equation, which is essential for a consistent convergence analysis based on a manufactured solution. This combination enables a more representative and rigorous assessment of the convergence of adaptive methods in this context.

To this end, we adopt a manufactured solution that defines a radially symmetric, time-dependent concentration field in d spatial dimensions:

$$\hat{u}(x,t) = \tanh\left[k\left(r^2(t) - \sum_{i=1}^{d} (x_i - x_{c\,i})^2\right)\right], \quad \boldsymbol{x} \in \mathbb{R}^{\mathbb{d}},$$

where the radius varies harmonically in time as

$$r(t) = r_0 + A\cos\left(\frac{\pi t}{T}\right).$$

For the simulations presented here, we focus on the 2D case $(d = 2)$ with the center of the computational domain and the phase-separated region given by $\mathbf{x}_c = (0.5, 0.5)$. The base radius and oscillation amplitude are set to $r_0 = 0.2$ and $A = 0.025$, respectively. The parameter $k = 100$ controls the sharpness of the interface, and $T = 3\,s$ denotes the oscillation period. This solution saturates at the binodal points ± 1, producing large, nearly uniform bulk phases separated by a smooth, time-evolving interface.

We consider a problem defined on the domain $\Omega = (0,1)^2$. In particular, we choose the source term $Q(\boldsymbol{x}, t)$ such that the exact solution is given by $\hat{u}(\boldsymbol{x}, t)$. To determine this source term, we substitute the exact solution into the strong form of the equation

$$\partial \hat{u}(\mathbf{x}, t)/\, \partial t - \nabla \cdot \left[\nabla \left(F'\big(\hat{u}(\mathbf{x}, t)\big) - \lambda \Delta \hat{u}(\mathbf{x}, t)\right)\right] = Q(\mathbf{x}, t)$$

where we assume constant and unit mobility. In the convergence studies performed for both the spatial and temporal discretisations, the source term $Q(\mathbf{x}, t)$ is introduced as an external body load and incorporated directly into the formulation. It is constructed from a manufactured solution so that the modified problem admits this solution exactly, enabling a controlled and systematic assessment of numerical errors. An illustration of the exact solution used is shown in Figure 3. Furthermore, since the analysis is carried out on a unit square domain, the essential boundary condition can be enforced strongly by assigning to the first layer of interior control points the value of the nearest boundary control point, following the approach of Kästner et al. [34].

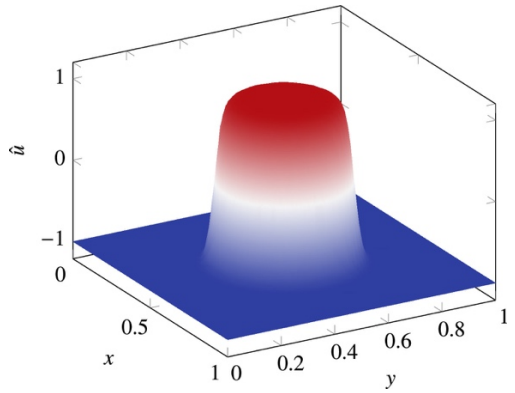

Figure 3.- Manufactured solution used for the convergence study, featuring a smooth, radially symmetric interface centered at $(x_c, y_c) = (0.5,0.5)$, shown at the initial time $t = 0\, s$.

Figure 4 shows the error evolution of the TP and THB discretizations using quadratic, cubic, and quartic basis functions. For quadratic basis functions ($p = 2$), presented in subfigures (a)-(c), THB discretizations achieve comparable accuracy to TP meshes with a reduction of roughly half an order of magnitude in degrees of freedom. This demonstrates that localized refinement can significantly improve computational efficiency without sacrificing solution accuracy.

In contrast, for cubic ($p = 3$) and quartic ($p = 4$) basis functions, shown in (d)-(f) and (g)-(i), respectively, THB meshes exhibit a loss of optimal convergence rates on finer meshes. This behavior is attributed to localized errors near the transition zones between the hierarchy levels within the phase-transition region and is observed primarily in the lower error norms. This highlights the critical role of the element marking strategy and underscores the inherent complexity of achieving low errors in convergence studies involving hierarchical meshes in the context of higher-order phase-field models. These observations underline the potential of THB discretizations to deliver accurate solutions at reduced computational cost; however, particular care must be taken when employing higher-degree basis functions, especially in the adaptive meshing process and in the construction of the manufactured solution. The latter poses a key limitation in this study, as the tanh-based manufactured solution does not perfectly saturate to the concentrations at the binodal points, which may lead to a loss of optimal convergence dominated by boundary condition enforcement errors.

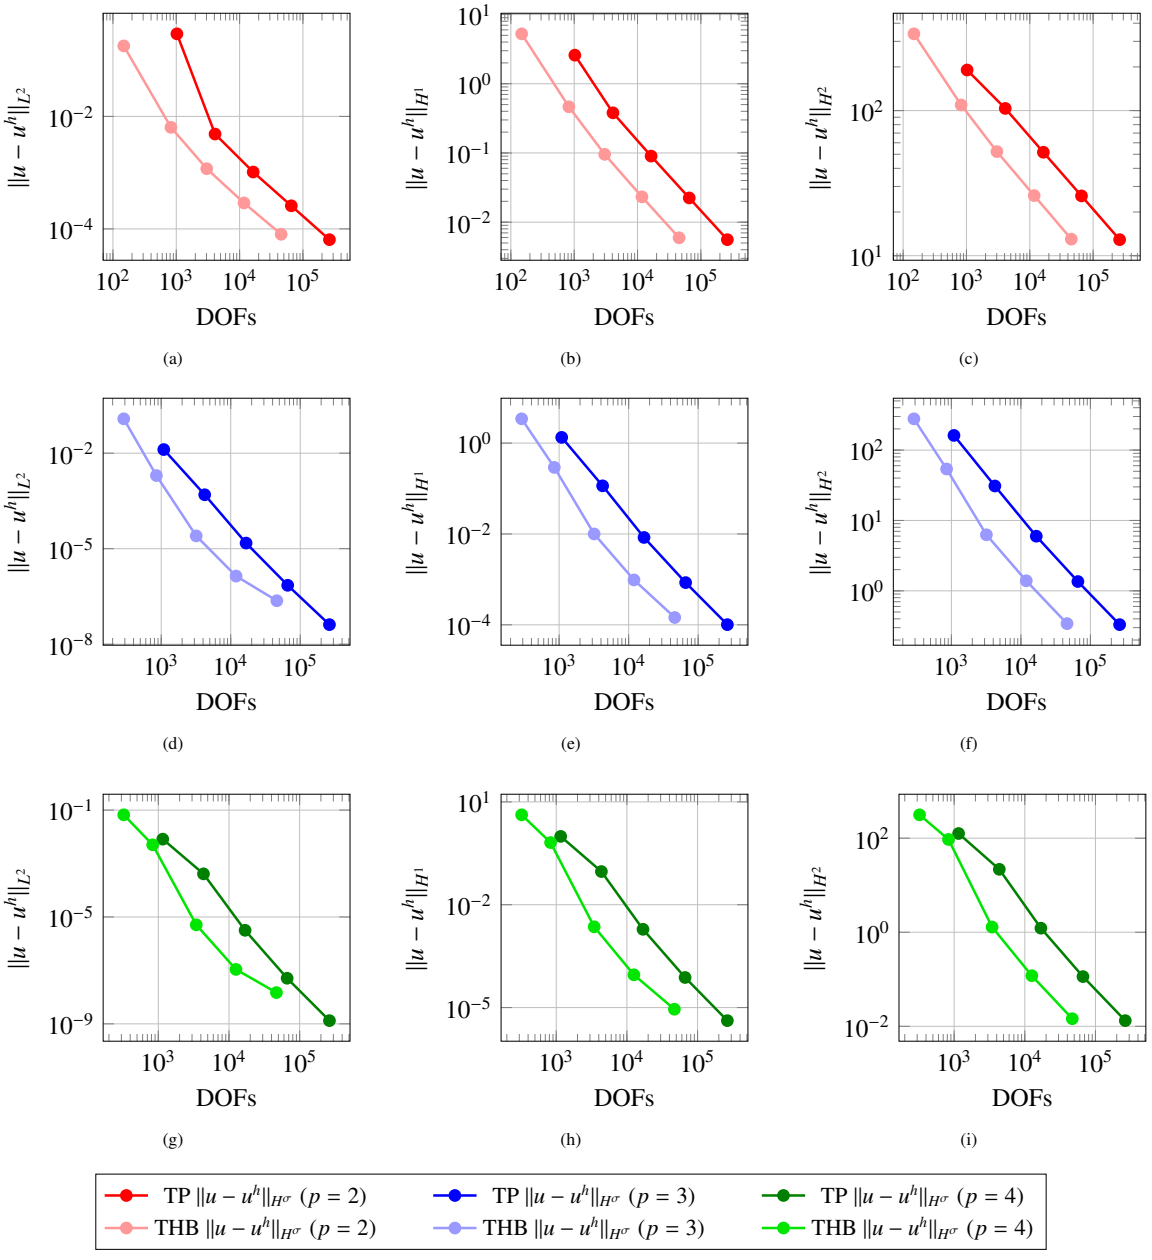

Figure 4.- Errors measured in the L^2-, H^1-, and H^2-norms, shown in (a)–(c) for quadratic ($p=2$), (d)–(f) for cubic ($p=3$), and (g)–(i) for quartic ($p=4$) basis functions, plotted against the number of degrees of freedom (DOFs) for both tensor product (TP) and truncated hierarchical B-spline (THB) meshes. Here, u denotes the manufactured (exact) solution and u^h the numerical solution.

4. Phase separation numerical examples

In this section, we investigate the performance of the adaptive isogeometric Cahn--Hilliard model. We limit our studies to simple geometrical domains to focus on the physical and numerical aspects of the problem. The domain of the test cases is a box $\Omega = [0,1]^d$, where $d = 2$ or 3. The essential boundary condition $\nabla u \cdot n = 0$ is enforced strongly. The spatial discretization is comprised of quadratic spline functions that are C^1-continuous across the knot spans, whereas time integration is based on the

generalized-α scheme. We employ meshes that are initially uniform in order to capture the rapid phase separation process that occurs at the beginning of the simulation with the formation of the interfaces.

The course of the phase separation process is influenced by the initial concentration distribution. Considering an average concentration $\tilde{u}$, different morphologies can be obtained [9]. Spinodal decomposition occurs when $\tilde{u}$ lies within the spinodal region, defined by the spinodal points of the double-well potential. In our formulation, this corresponds to $\tilde{u} = 0.0$ and results in interconnected, bicontinuous morphologies. Conversely, when $\tilde{u}$ lies in the metastable region —between the binodal and spinodal curves— one phase is thermodynamically favored, and phase separation proceeds via nucleation-and-growth, producing isolated domains rather than a bicontinuous morphology.

In all cases, the simulations are initialized from a perturbed state

$$u(x, 0) = \tilde{u} + \delta u(x),$$

where $\delta u(x)$ is a uniformly distributed random perturbation.

In all simulations, the finest level of the domain of interest Ω consists of 64^d uniformly distributed quadratic elements, where d represents the problem dimension. In the case of adaptive simulations, elements of coarser levels are activated by the adaptive meshing algorithm presented in this paper, depending on the concentration u, the mesh configuration, and the refinement/coarsening threshold parameter $\Theta = 0.1$. Admissible mesh refinement is obtained with grading parameter $m = 2$. The following model parameters are assumed: $M = 1\,\mathrm{m^3\,s\,kg^{-1}}, \lambda = 1.5 \times 10^{-3}\,\mathrm{J\,m^{-1}}$. We adopt a constant time step of $(\Delta t = 1.5 \times 10^{-3}\sim\mathrm{s})$ over a total simulation time of $(T = 1.5\mathrm{s})$.

All numerical experiments were performed using the *G+Smo* library [39] on an Intel Xeon Silver 4410Y CPU. Linear systems were assembled using shared-memory parallelization and solved with the PARDISO solver from the Intel Math Kernel Library, using 40 threads of the allocated CPUs. The quasi-interpolation approximants were also constructed using shared-memory parallelization with the same number of threads. Reported computational times correspond to this computational environment.

4.1. Nucleation-driven phase separation in 2D

Figure 5 illustrates the temporal evolution of the concentration field starting from an initial average concentration of $\tilde{u} = 0.4$ with a small random perturbation $\delta u(x) \in [-0.005, 0.005]$, leading to the formation of circular inclusions in the steady state. Three representative time steps are shown: the first row displays the TP mesh solution, while the second and third rows show adaptive solutions obtained with L^2-projection and quasi-interpolant-based transfer, respectively.

Figure 6 summarizes the nucleation results. Subfigure (a) shows that the adaptive schemes reduce DOFs by roughly a factor of two compared to the TP solution. Subfigure (b) confirms that all simulations follow the same energetic evolution, while subfigure (c) demonstrates excellent mass conservation. Computational performance per refinement iteration is reported in subfigures (d) to (f), showing the assembly, solver, and projection times, respectively. Notably, the quasi-interpolant reduces the projection time by approximately 65\% compared to the L^2-projection.

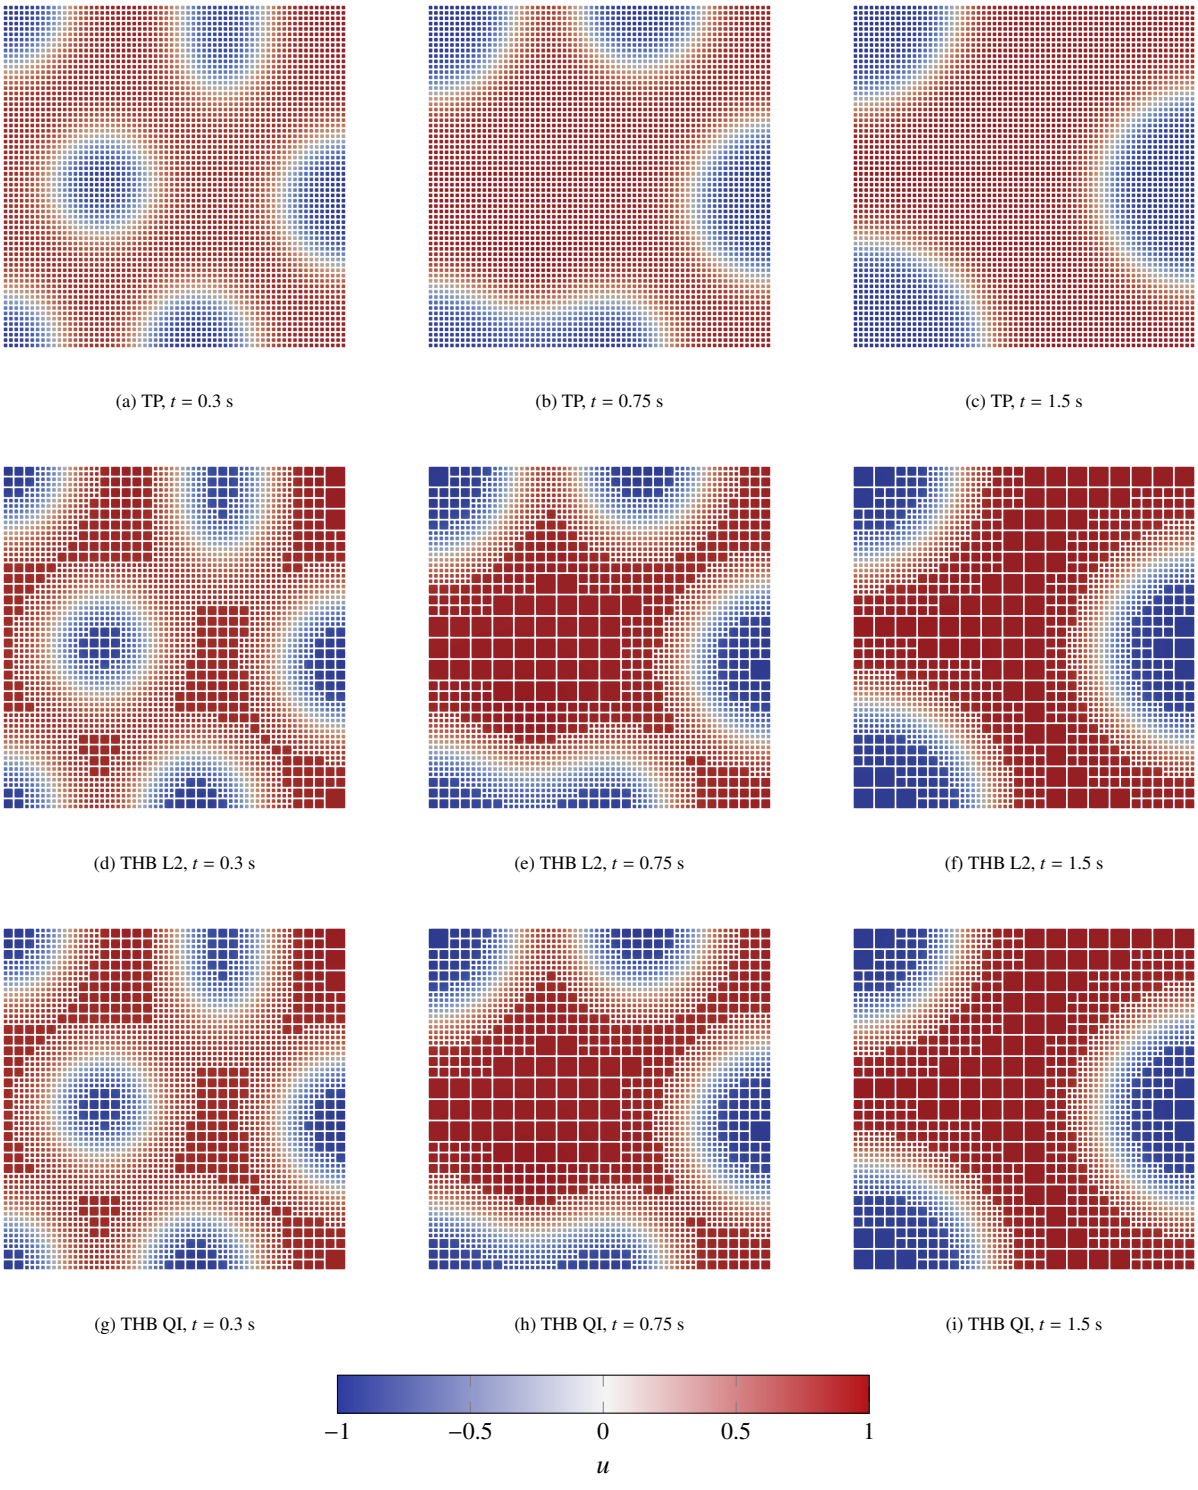

Figure 5.- Nucleation in 2D. Evolution of the concentration from an initially perturbed state $u(x,0) = \tilde{u} + \delta u(x)$, with $\tilde{u} = 0.4$ and $\delta u(x)$ a small random perturbation uniformly distributed in $[-0.005, 0.005]$. Top row (a-c): tensor product (TP) mesh; middle row (d-f): adaptivity using L^2-projection (THB L2); bottom row (g-i): adaptivity using the quasi-interpolant based on local interpolation (THB QI). The mesh consists of 64^2 quadratic elements in the TP construction.

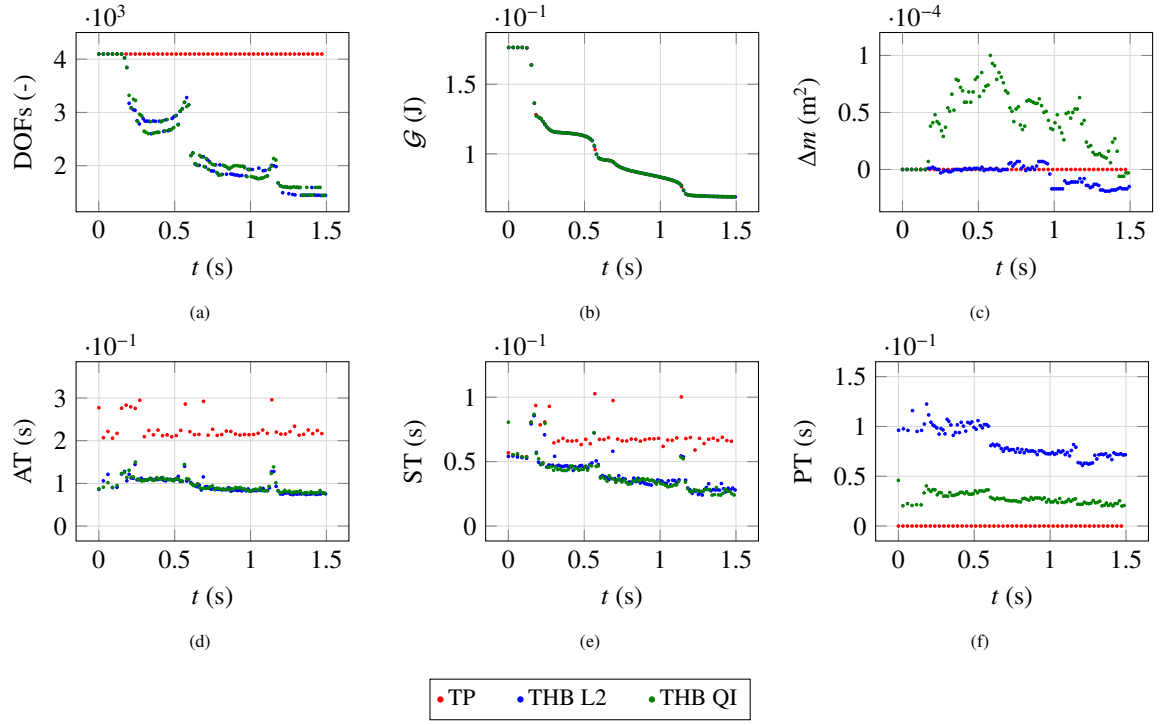

Figure 6.- Overview of the computational insights during nucleation in 2D, comparing three approaches — tensor product (TP), adaptivity with L^2-projection (THB L2), and adaptivity with quasi-interpolant (THB QI) — for the evolution of degrees of freedom (DOFs) (a), total energy $\mathcal{G}$(b), and mass difference $\Delta m = \int_\Omega u \, d\Omega - \int_\Omega u_{\mathrm{TP}} \, d\Omega$ (c), where u_{TP} is the solution computed with the TP mesh. Computational times are reported for assembly time (AT, d), solver time (ST, e), and projection time (PT, f). The initial condition is $u(x,0) = u_0 + \delta u(x)$ with $u_0 = 0.4$ and $\delta u(x)$ a small random perturbation uniformly distributed in $[-0.005, 0.005]$. The mesh consists of 64^2 quadratic elements in the TP construction.

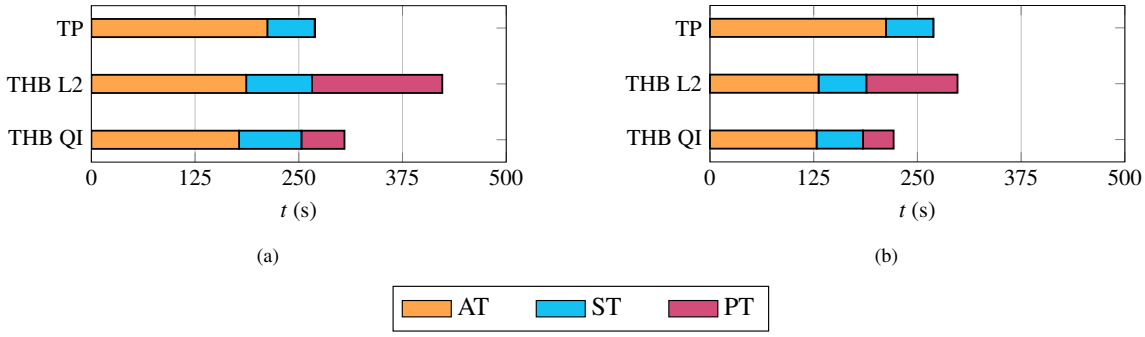

Figure 7.- Total computational times for 2D nucleation using three approaches: tensor product (TP), adaptivity with L^2-projection (THB L2), and adaptivity with the quasi-interpolant (THB QI). Two refinement strategies are compared: (a) simulations without a maximum number of refinement iterations, and (b) simulations where refinement iterations are triggered by significant changes in the total energy.

Figure 7 compares the total simulation times for the three cases, highlighting the overall computational cost. The difference between the two strategies lies in how the refinement loop is activated. In the first approach, the loop is activated exactly as prescribed in Figure 1, whenever at least one element is marked for refinement, i.e., $|\mathcal{M}_r| > 0$. This causes additional assembly, solution, and projection steps even for minor discretization changes, thereby introducing substantial computational overhead (see

subfigure (a)). In the second approach, the refinement loop is executed only during selected step intervals chosen based on observed system evolution, which approximately corresponds to periods of significant energy change. This manual selection reduces the number of refinement iterations and, consequently, the total computational cost (see subfigure (b)), resulting in a marked reduction in total simulation time. Specifically, the total simulation time decreases by approximately 18% for the L^2-projection and 12% for the quasi-interpolant.

5. Conclusions and outlook

This report has presented the development of an adaptive isogeometric framework for higher-order phase-field modeling, focusing on the Cahn–Hilliard equation as a representative problem for interfacial dynamics. The results demonstrate that the combination of smooth spline-based discretizations and hierarchical adaptive refinement strategies provides an accurate and computationally efficient approach for simulating phase separation processes. The high-order continuity of isogeometric basis functions enables a direct primal formulation without the need for auxiliary variables. In addition, THB-splines facilitate localized refinement and coarsening, allowing computational resources to be focused on regions of interest, such as evolving interfaces.

The numerical investigations confirm the robustness and reliability of the proposed methodology. Convergence studies validate the suitability of hierarchical spline spaces for higher-order problems, and simulations with complex initial conditions highlight the ability of the adaptive scheme to accurately capture interface evolution. At the same time, the reduction in degrees of freedom compared to uniform meshes illustrates a favorable balance between accuracy and computational cost, making the approach particularly attractive for large-scale simulations.

Overall, this work contributes to advancing adaptive phase-field modeling by providing an efficient and flexible computational framework. Beyond phase separation, the methodology shows strong potential for applications in other interfacial problems where localized features dominate the solution behavior.

Future research will focus on extending the framework to fully three-dimensional and more complex geometries, as well as to advanced phase-field models such as fracture and multiphysics systems. In parallel, efforts will be devoted to improving computational performance through optimized assembly techniques, including sum-factorization and weighted quadrature on hierarchical meshes, and the development of scalable solvers for large linear systems. These aspects are particularly critical in 3D simulations, where solver efficiency becomes dominant. Finally, the exploration of adaptive immersed methods will further enhance the flexibility of the approach, enabling the treatment of complex geometries without boundary-fitted meshes and strengthening the integration between geometric design and numerical analysis.

6. References

[1] J. W. Cahn and J. E. Hilliard, “Free Energy of a Nonuniform System. I. Interfacial Free Energy,” *J. Chem. Phys.*, vol. 28, no. 2, pp. 258–267, Feb. 1958, doi: 10.1063/1.1744102.
[2] J. W. Cahn and J. E. Hilliard, “Free Energy of a Nonuniform System. III. Nucleation in a Two-Component Incompressible Fluid,” *J. Chem. Phys.*, vol. 31, no. 3, pp. 688–699, Sep. 1959, doi: 10.1063/1.1730447.
[3] P. G. Ciarlet and P. A. Raviart, “A Mixed Finite Element Method for the Biharmonic Equation,” in *Mathematical Aspects of Finite Elements in Partial Differential Equations*, Elsevier, 1974, pp. 125–145. doi: 10.1016/B978-0-12-208350-1.50009-1.
[4] D. Furihata, “A stable and conservative finite difference scheme for the Cahn-Hilliard equation:,” *Numer. Math.*, vol. 87, no. 4, pp. 675–699, Feb. 2001, doi: 10.1007/PL00005429.
[5] G. N. Wells, E. Kuhl, and K. Garikipati, “A discontinuous Galerkin method for the Cahn–Hilliard equation,” *J. Comput. Phys.*, vol. 218, no. 2, pp. 860–877, Nov. 2006, doi: 10.1016/j.jcp.2006.03.010.
[6] T. J. R. Hughes, J. A. Cottrell, and Y. Bazilevs, “Isogeometric analysis: CAD, finite elements, NURBS, exact geometry and mesh refinement,” *Comput. Methods Appl. Mech. Eng.*, vol. 194, no. 39–41, pp. 4135–4195, Oct. 2005, doi: 10.1016/j.cma.2004.10.008.
[7] J. Kiendl, K.-U. Bletzinger, J. Linhard, and R. Wüchner, “Isogeometric shell analysis with Kirchhoff–Love elements,” *Comput. Methods Appl. Mech. Eng.*, vol. 198, no. 49–52, pp. 3902–3914, Nov. 2009, doi: 10.1016/j.cma.2009.08.013.
[8] J. Kiendl, M.-C. Hsu, M. C. H. Wu, and A. Reali, “Isogeometric Kirchhoff–Love shell formulations for general hyperelastic materials,” *Comput. Methods Appl. Mech. Eng.*, vol. 291, pp. 280–303, Jul. 2015, doi: 10.1016/j.cma.2015.03.010.
[9] H. Gómez, V. M. Calo, Y. Bazilevs, and T. J. R. Hughes, “Isogeometric analysis of the Cahn–Hilliard phase-field model,” *Comput. Methods Appl. Mech. Eng.*, vol. 197, no. 49–50, pp. 4333–4352, Sep. 2008, doi: 10.1016/j.cma.2008.05.003.
[10] M. J. Borden, C. V. Verhoosel, M. A. Scott, T. J. R. Hughes, and C. M. Landis, “A phase-field description of dynamic brittle fracture,” *Comput. Methods Appl. Mech. Eng.*, vol. 217–220, pp. 77–95, Apr. 2012, doi: 10.1016/j.cma.2012.01.008.
[11] G. Lorenzo, T. J. R. Hughes, P. Dominguez-Frojan, A. Reali, and H. Gomez, “Computer simulations suggest that prostate enlargement due to benign prostatic hyperplasia mechanically impedes prostate cancer growth,” *Proc. Natl. Acad. Sci.*, vol. 116, no. 4, pp. 1152–1161, Jan. 2019, doi: 10.1073/pnas.1815735116.
[12] L. Greco, E. Maggiorelli, M. Negri, A. Patton, and A. Reali, “AT1 fourth-order isogeometric phase-field modeling of brittle fracture,” *Math. Models Methods Appl. Sci.*, pp. 1–55, Sep. 2025, doi: 10.1142/S0218202525500502.
[13] Y. Bazilevs, V. M. Calo, T. J. R. Hughes, and Y. Zhang, “Isogeometric fluid-structure interaction: theory, algorithms, and computations,” *Comput. Mech.*, vol. 43, no. 1, pp. 3–37, Dec. 2008, doi: 10.1007/s00466-008-0315-x.
[14] Y. Bazilevs, C. Michler, V. M. Calo, and T. J. R. Hughes, “Isogeometric variational multiscale modeling of wall-bounded turbulent flows with weakly enforced boundary conditions on unstretched meshes,” *Comput. Methods Appl. Mech. Eng.*, vol. 199, no. 13–16, pp. 780–790, Feb. 2010, doi: 10.1016/j.cma.2008.11.020.
[15] C. Bracco, C. Giannelli, A. Reali, M. Torre, and R. Vázquez, “Adaptive isogeometric phase-field modeling of the Cahn–Hilliard equation: Suitably graded hierarchical refinement and coarsening on multi-patch geometries,” *Comput. Methods Appl. Mech. Eng.*, vol. 417, p. 116355, Dec. 2023, doi: 10.1016/j.cma.2023.116355.
[16] A. Buffa, G. Gantner, C. Giannelli, D. Praetorius, and R. Vázquez, “Mathematical Foundations of Adaptive Isogeometric Analysis,” *Arch. Comput. Methods Eng.*, vol. 29, no. 7, pp. 4479–4555, Nov. 2022, doi: 10.1007/s11831-022-09752-5.
[17] C. De Boor, *A practical guide to splines*. in Applied mathematical sciences, no. 27. Berlin New York Paris [etc.]: Springer, 1978.

[18] D. R. Forsey and R. H. Bartels, “Hierarchical B-spline refinement,” *ACM SIGGRAPH Comput. Graph.*, vol. 22, no. 4, pp. 205–212, Aug. 1988, doi: 10.1145/378456.378512.
[19] C. Giannelli, B. Jüttler, and H. Speleers, “THB-splines: The truncated basis for hierarchical splines,” *Comput. Aided Geom. Des.*, vol. 29, no. 7, pp. 485–498, Oct. 2012, doi: 10.1016/j.cagd.2012.03.025.
[20] A. Buffa and C. Giannelli, “Adaptive isogeometric methods with hierarchical splines: Error estimator and convergence,” *Math. Models Methods Appl. Sci.*, vol. 26, no. 01, pp. 1–25, Jan. 2016, doi: 10.1142/S0218202516500019.
[21] X. Li and M. A. Scott, “Analysis-suitable T-splines: Characterization, refineability, and approximation,” *Math. Models Methods Appl. Sci.*, vol. 24, no. 06, pp. 1141–1164, Jun. 2014, doi: 10.1142/S0218202513500796.
[22] X. Li and J. Zhang, “AS++ T-splines: Linear independence and approximation,” *Comput. Methods Appl. Mech. Eng.*, vol. 333, pp. 462–474, May 2018, doi: 10.1016/j.cma.2018.01.041.
[23] X. Li and T. W. Sederberg, “S-splines: A simple surface solution for IGA and CAD,” *Comput. Methods Appl. Mech. Eng.*, vol. 350, pp. 664–678, Jun. 2019, doi: 10.1016/j.cma.2019.03.035.
[24] N. Nguyen-Thanh, H. Nguyen-Xuan, S. P. A. Bordas, and T. Rabczuk, “Isogeometric analysis using polynomial splines over hierarchical T-meshes for two-dimensional elastic solids,” *Comput. Methods Appl. Mech. Eng.*, vol. 200, no. 21–22, pp. 1892–1908, May 2011, doi: 10.1016/j.cma.2011.01.018.
[25] T. Dokken, T. Lyche, and K. F. Pettersen, “Polynomial splines over locally refined box-partitions,” *Comput. Aided Geom. Des.*, vol. 30, no. 3, pp. 331–356, Mar. 2013, doi: 10.1016/j.cagd.2012.12.005.
[26] E. M. Garau and R. Vázquez, “Algorithms for the implementation of adaptive isogeometric methods using hierarchical B-splines,” *Appl. Numer. Math.*, vol. 123, pp. 58–87, Jan. 2018, doi: 10.1016/j.apnum.2017.08.006.
[27] S. Nagaraja, M. Elhaddad, M. Ambati, S. Kollmannsberger, L. De Lorenzis, and E. Rank, “Phase-field modeling of brittle fracture with multi-level hp-FEM and the finite cell method,” *Comput. Mech.*, vol. 63, no. 6, pp. 1283–1300, Jun. 2019, doi: 10.1007/s00466-018-1649-7.
[28] L. Chen, B. Li, and R. De Borst, “Adaptive isogeometric analysis for phase-field modeling of anisotropic brittle fracture,” *Int. J. Numer. Methods Eng.*, vol. 121, no. 20, pp. 4630–4648, Oct. 2020, doi: 10.1002/nme.6457.
[29] H. Speleers and C. Manni, “Effortless quasi-interpolation in hierarchical spaces,” *Numer. Math.*, vol. 132, no. 1, pp. 155–184, Jan. 2016, doi: 10.1007/s00211-015-0711-z.
[30] A. Giust, B. Jüttler, and A. Mantzaflaris, “Local (T)HB-spline projectors via restricted hierarchical spline fitting,” *Comput. Aided Geom. Des.*, vol. 80, p. 101865, Jun. 2020, doi: 10.1016/j.cagd.2020.101865.
[31] J. A. Evans, Y. Bazilevs, I. Babuška, and T. J. R. Hughes, “n-Widths, sup–infs, and optimality ratios for the k-version of the isogeometric finite element method,” *Comput. Methods Appl. Mech. Eng.*, vol. 198, no. 21–26, pp. 1726–1741, May 2009, doi: 10.1016/j.cma.2009.01.021.
[32] M. L. Huggins, “Solutions of Long Chain Compounds,” *J. Chem. Phys.*, vol. 9, no. 5, pp. 440–440, May 1941, doi: 10.1063/1.1750930.
[33] P. J. Flory, “Thermodynamics of High Polymer Solutions,” *J. Chem. Phys.*, vol. 10, no. 1, pp. 51–61, Jan. 1942, doi: 10.1063/1.1723621.
[34] M. Kästner, P. Metsch, and R. De Borst, “Isogeometric analysis of the Cahn–Hilliard equation – a convergence study,” *J. Comput. Phys.*, vol. 305, pp. 360–371, Jan. 2016, doi: 10.1016/j.jcp.2015.10.047.
[35] M. Carraturo, C. Giannelli, A. Reali, and R. Vázquez, “Suitably graded THB-spline refinement and coarsening: Towards an adaptive isogeometric analysis of additive manufacturing processes,” *Comput. Methods Appl. Mech. Eng.*, vol. 348, pp. 660–679, May 2019, doi: 10.1016/j.cma.2019.01.044.
[36] H. M. Verhelst, A. Mantzaflaris, M. Möller, and J. H. Den Besten, “Goal-adaptive Meshing of Isogeometric Kirchhoff–Love Shells,” *Eng. Comput.*, vol. 40, no. 6, pp. 3595–3622, Dec. 2024, doi: 10.1007/s00366-024-01958-4.
[37] W. Dörfler, “A Convergent Adaptive Algorithm for Poisson’s Equation,” *SIAM J. Numer. Anal.*, vol. 33, no. 3, pp. 1106–1124, Jun. 1996, doi: 10.1137/0733054.

[38] A. Tagliabue, L. Dedè, and A. Quarteroni, “Isogeometric Analysis and error estimates for high order partial differential equations in fluid dynamics,” *Comput. Fluids*, vol. 102, pp. 277–303, Oct. 2014, doi: 10.1016/j.compfluid.2014.07.002.
[39] B. Jüttler, U. Langer, A. Mantzaflaris, S. E. Moore, and W. Zulehner, “Geometry + Simulation Modules: Implementing Isogeometric Analysis,” *PAMM*, vol. 14, no. 1, pp. 961–962, Dec. 2014, doi: 10.1002/pamm.201410461.