



Semi-intrusive stiffener path optimisation for hybrid manufacturing

Duarte F. Valoroso Madeira^{1,2} · Armin Geiser² · Majid Hojjat² · Benedikt Kriegesmann¹

Received: 13 April 2026 / Revised: 5 August 2026 / Accepted: 27 August 2026
© The Author(s) 2026

Abstract

The paper presents a new semi-intrusive method to optimise the paths of stiffeners on complex surfaces. One motivation for such approaches is hybrid manufacturing, where a conventionally manufactured shell structure is reinforced by additively manufactured stiffeners. This potential application motivates manufacturing constraints, which are embedded in the new approach. These are, in particular, the requirements that stiffeners should be orthogonal to the geometry they are being printed on, and that parallel stiffeners should have a minimum distance, whilst it shall be allowed for stiffeners to cross. The main challenge when optimising stiffeners on a shell is that connectivity of stiffeners and shell must be maintained, but remeshing of the shell is not desirable. Coupling degrees of freedom via element shape functions however requires full access to the finite element code. In our approach we circumvent both by performing only local mesh modifications. The aforementioned manufacturing constraints are implemented as a layered normal-based parametrisation and an angle-based penalised distance constraint. Everything is implemented in a projection-based setting that allows the stiffeners path to vary along an arbitrarily shaped surface. The applicability of the framework is demonstrated by academic benchmark examples as well as a more challenging industrial-like example.

Keywords Semi-intrusive method · Non-conforming mesh · Shape optimisation · Stiffened structures · Hybrid manufacturing

1 Introduction

Additive manufacturing is always associated with design freedom. In hybrid manufacturing (i.e. combined additive and conventional manufacturing), one notable use case for structural applications involves using additive manufacturing to fabricate stiffeners (Kloft et al. 2023; Borg Costanzi 2023). Wire and Arc Additive Manufacturing (WAAM) is leveraged to print stiffeners onto base geometries (cf. Fig. 1), which mainly consist of already manufactured sheet-metal components that are dimensionally fit for existing assemblies. Adding stiffeners enables increasing the stiffness for targeted use cases without redesigning the part, thus reducing lead times and cost. Hence, the goal is to optimise stiffeners based on existing geometries that previously did not meet the mechanical and structural requirements.

Stiffener optimisation is typically done in one of two ways, using a continuous design space via topology optimisation, or discretely modelling stiffeners with shape optimisation. Whilst a continuous design space offers broad design freedom and smooth parametrisations such as B-spline height fields on curved shells with solid-shell coupling (Feng et al. 2021) and anisotropic Helmholtz filtering to embed casting constraints (Zhou et al. 2022), it faces significant challenges. Achieving adequate resolution for thin, high-aspect-ratio stiffeners at low volume fractions demands extremely fine meshes, often involving millions or even billions of voxels. This leads to excessive computational costs, limiting applicability mostly to academic-scale or smaller components. Even with giga-voxel SIMP morphogenesis enabled by super-computing resources (Aage et al. 2017), the resulting geometries tend to be difficult to manufacture due to intermediate density regions. Furthermore, continuous density-based methods struggle to reliably produce designs with volume fractions below ~20 %, where discrete stiffener formulations have been shown to provide superior stiffness and more robust manufacturability (Carstensen et al. 2023).

Using discrete stiffeners we allow their paths to vary continuously, but persistent trade-offs between design freedom,

Responsible Editor: Graeme James Kennedy.

✉ Duarte F. Valoroso Madeira
duarte.madeira@tuhh.de; duarte.madeira@bmw.de

¹ Institute SML, Hamburg University of Technology,
Eißendorfer Straße 40 (N), 21073 Hamburg, Germany

² BMW Group, Petuelring 130, Munich, Germany

computational tractability, and manufacturability characterise all approaches in this category. Orientation optimisation of isogrid plates has demonstrated measurable performance gains with compact FE models and experimental validation (Akl et al. 2008). However, design freedom concentrates on stiffener angles and pattern regularity, limiting applicability to free-form layouts and complex curvature mappings. Sensitivity-driven reinforcement placement using adjoint evaluations can efficiently localise material in 2D plate contexts (Bojczuk and Szeleblak 2008), yet requires highly intrusive coupling between design and analysis. Hence, differentiating the coupled base–stiffener system is not readily supported in commercial FEA environments, restricting practical implementation. Geometry-exact frameworks based on level sets and IGA have coupled topology, shape, and sizing with smooth analysis, e.g. embedded Kirchhoff–Love shells with NURBS composition and mortar coupling (Hirschler et al. 2019), and level-set stiffened panels with p-norm stress aggregation and multi-mode buckling constraints (Chu et al. 2021), and an isogeometric design analysis optimisation workflow that couples non-conforming NURBS patches to optimise skin shape and stiffener layout concurrently (Hao et al. 2023), but these methods restrict topological operations such as stiffener crossings, remain highly intrusive, are, in some cases, dependent on a dedicated solver with the stiffener topology fixed by the patch layout, and face challenges in ensuring IGA continuity and stability, handling seamlines, and navigating bounded search spaces induced by fixed patterns. A surrogate-based scheme for curvilinear, gradual-stiffness panels defined on complex NURBS surfaces (Hao et al. 2024) avoids the intrusiveness of the above methods, as it requires no sensitivity information and analyses the design with commercial FE software, and relaxes the fixed two-cluster orthogonal stiffener pattern via a stiffener-unit-cell database, but its dimension-reduction metamodeling instead limits the number of design variables it can handle. Mulani et al. (2013) developed a PSO-driven optimisation of curvilinear stiffeners within a commercial FEA framework, but it suffers from poor scalability as its heuristic optimisation approach struggles with dimensionality, making it less effective than gradient-based methods (Sigmund 2011).

Finally, other approaches for this problem use a ground-structure mesh, a foundational, highly connected network of potential beam elements used in structural optimisation to identify the most efficient load-bearing topology. Approaches combining both topology and shape optimisation have been developed on this method, where coupled plate–beam element models enable simultaneous topology, shape and sizing optimisation of ribbed plate systems (Amir and Majdoub 2024). An IGA-based stiffness spreading method has also been explored to achieve smooth distributions (Sun et al. 2022). Zhao et al. (2022) presented a




Fig. 1 Stiffened sheet-metal geometry

two-level approximation method combining genetic algorithms for topology optimisation and gradient-based shape and size refinement, effectively mixing discrete and continuous variables, which complicates solver integration and may affect convergence. The major issue with all these approaches is the inherent limitations of using a ground mesh. The extensive amount of connectivities between nodes results in a very dense stiffness matrix and consequently compromising scalability of the problem. Furthermore, since design freedom is constrained, optimal layouts may be missed, especially since mesh distortion needs to be limited, requiring significantly more conservative step lengths (Weldeyesus et al. 2020) and relying more on a good initial layout guess.

In summary, the reviewed methods share several persistent challenges: (i) limited scalability, (ii) difficulties in handling arbitrarily curved shells and complex geometries, and (iii) the absence of WAAM-specific manufacturing constraints.

To address these challenges, this work introduces a semi-intrusive gradient-based framework specifically designed for industrial deployment. The proposed method employs a local element subdivision scheme that couples non-conforming stiffener and base meshes without requiring access to proprietary shape function formulations, enabling seamless integration with commercial FEA software and scalability to industrial-sized parts. This coupling removes much of the intrusiveness of a classical implementation, whilst the design is updated through a gradient-based optimiser rather than a black-box heuristic. The approach therefore lies between a fully non-intrusive plug-in and a fully intrusive formulation, which is what the term semi-intrusive denotes. The shape optimisation is driven by vertex morphing (Hojjat et al. 2014; Bletzinger 2014) which enables smooth, controllable geometric updates unconstrained by predefined candidate networks. Unlike existing approaches limited to flat or simple curved geometries, the framework handles arbitrarily curved shells through in-plane projection, allowing stiffener paths to follow complex surface curvatures typical of real automotive and aerospace components. The framework embeds manufacturing-specific features tailored to WAAM through two key mechanisms: a normal-based parametrisation

tion that ensures stiffener orthogonality with regards to the base geometry and an angle-based distance constraint that permits stiffener crossings whilst enforcing minimum distance only for parallel or near-parallel members.

The paper is structured as follows. Section 2 recapitulates the theoretical background required for the new approach and its validation. Section 3 presents the new semi-intrusive approach for stiffener path optimisation. Section 4 presents the manufacturing constraints. Numerical results of different examples are summarised in Sect. 5.

2 Theoretical background and problem formulation

This section presents the theoretical foundations underlying the proposed framework. First, we formulate the constrained shape optimisation problem and introduce the optimisation algorithms used to solve it. Subsequently, we describe the vertex morphing parametrisation and its extension for curved geometries through in-plane projection. Finally, we present the intrusive Lagrange multiplier coupling approach, which establishes the mathematical basis for connecting non-conforming meshes and serves as the reference for validation in Sect. 3.

2.1 Constrained shape optimisation

Specific to the present stiffener path optimisation, the constrained shape optimisation problem, targeting the minimisation of a structural response function subject to geometric and manufacturing constraints, is formulated as:

$$\begin{aligned} \min \quad & f(\mathbf{x}) \\ \text{s.t.} \quad & p(\mathbf{x}) \leq 0, \\ & d(\mathbf{x}) \geq d_{\min}, \\ & m(\mathbf{x}) \leq m_{\max}. \end{aligned} \quad (1)$$

where $f(\mathbf{x})$ is the objective function, here the structural compliance, and $\mathbf{x}$ are the design nodes, which are the only design variables. The constraints express the geometric and manufacturing requirements of this specific problem, with their detailed formulation deferred to the corresponding sections: the packaging constraint $p(\mathbf{x})$ is a no-penetration constraint following (Asl et al. 2017), evaluated point-wise as a gap function that measures the signed normal distance from each design node to its closest point on a prescribed bounding surface, here a packaging box, so that enforcing $p(\mathbf{x}) \leq 0$ keeps the design nodes on the feasible side of that surface. The distance constraint $d(\mathbf{x})$ enforces a minimum spacing $d_{\min}$ between neighbouring stiffeners, described later in Sect. 4.2.

Finally, $m(\mathbf{x})$ is a simple mass response that keeps it below a prescribed limit $m_{\max}$.

To solve this constrained problem, one can use the Relaxed Gradient Projection (RGP) method, which projects the steepest descent direction onto the feasible region defined by active constraints (Antonau et al. 2021). However, RGP tends to produce an oscillatory shape update behaviour, particularly when design variables approach constraint boundaries where nodes repeatedly violate and correct said constraints. This consequently results in a very noisy optimisation.

The Quasi-Newton Barzilai-Borwein method computes each design variable's step size independently based on the local sensitivity information. Coupling this with the RGP method algorithm yields the Quasi-Newton Relaxed Gradient Projection (QNRGP) method (Antonau et al. 2022) that effectively decelerates nodal updates as they approach constraint boundaries, providing smoother convergence behaviour. Rather than abrupt corrections at constraint violations, the algorithm anticipates proximity to constraints and gradually reduces step sizes, resulting in shapes that respect packaging limits. Given the large number of geometric and manufacturing constraints involved in the present stiffener path optimisation, the QNRGP method is the one adopted throughout this work, precisely for its more robust handling of active constraint boundaries, which avoids the noisy updates produced by the plain RGP method.

2.2 Vertex morphing

Vertex morphing (Hojjat et al. 2014; Bletzinger 2014) introduces a design control field s that controls the shape x via convolution with a filter function F , i.e. a kernel function, here taken as the linear hat function:

$$x(\xi_i) = \int_{\xi_i-r}^{\xi_i+r} F(\xi_i, \xi, r) s \, d\xi \quad (2)$$

where ξ_i is the surface coordinate of the unknown point and r is the filter radius, which governs the smoothness of the generated shape. The filter radius can be seen as an additional design parameter controlling the desired shape modes of the optimal design. As the method operates in 3D, r is a spherical radius, kept constant throughout the optimisation.

As shape smoothness is set mainly by the filter radius rather than the kernel, smoother kernels could be used with little effect (Bletzinger 2014). Moreover, the hat function vanishes at its support boundaries, so nodes moving in and out of the filter contribute negligibly. Therefore, this small parametrisation non-linearity is neglected since the shape updates remain smooth regardless.

In a discrete setting, the shape and the control field are represented by the nodal vectors $\mathbf{x}$ and $\mathbf{s}$, respectively. Adopting the same discretisation for both fields, the convolution in

Eq. (2) reduces to a linear map described by a filter matrix $\mathbf{F}$, whose entries collect the kernel weights $F(\xi_i, \xi, r)$. This map is applied twice in every optimisation cycle. First, a control field update $\delta \mathbf{s}$ is forward filtered to generate the free-form shape update

$$\delta \mathbf{x} = \mathbf{F} \delta \mathbf{s}, \quad (3)$$

and second, the response function gradients evaluated on the geometry are backward filtered onto the control field through the transpose of the same matrix,

$$\frac{\partial J}{\partial \mathbf{s}} = \mathbf{F}^T \frac{\partial J}{\partial \mathbf{x}}. \quad (4)$$

The forward filter therefore yields precisely the free-form shape update $\delta \mathbf{x}_i$ that the in-plane projection subsequently modifies, whereas the backward filter provides the gradient mapping invoked in Sect. 2.2.1. Together, the forward and backward filtering establish the link between the vertex morphing parametrisation and its projection onto an arbitrarily curved reference surface.

2.2.1 In-plane projection

A key challenge in applying shape optimisation to industrial components, particularly in the automotive sector, is handling geometries with complex, irregular curvatures. Unlike academic benchmarks that often consider flat plates or simple curved shells, real-world parts exhibit non-trivial surface topologies where stiffeners must follow the underlying geometry whilst optimising their paths. To address this, we use an in-plane projection, which forces all surface points to slide along a reference surface rather than moving freely in 3D space.

The in-plane projection operates in two steps. First, the free-form shape update from Eq. (3) is projected onto the tangent plane of the reference surface at the current position, defined by the surface normal $\mathbf{n}_i$. In a second step, any violation of the sliding condition due to surface curvature is corrected:

$$\delta \tilde{\mathbf{x}}_i = \delta \mathbf{x}_i - \delta \mathbf{x}_i \mathbf{n}_i \cdot \mathbf{n}_i + \mathbf{c}(\delta \mathbf{x}_i) \quad (5)$$

where $\mathbf{c}$ is the curvature correction term. In this way, Eq. (5) extends the forward filter of Eq. (3): rather than being applied directly, the free-form update $\delta \mathbf{x}_i$ is projected onto the tangent plane and corrected for curvature to yield the in-plane update $\delta \tilde{\mathbf{x}}_i$. Following the chain rule of differentiation, a mapping for the response function gradients can be determined,




Fig. 2 Modification of the free-form shape update by projection into the tangent plane and correction due to curvature

which extends the backward filter of Eq. (4) in the same manner, so that the projected and curvature-corrected sensitivities are the ones propagated back onto the control field. The correction step due to surface curvature introduces a non-linear relation between the parametrisation and the shape update. A visual representation of the projection and correction process is illustrated in Fig. 2.

In contrast to the linear relation in standard vertex morphing, this non-linearity cannot be accounted for in the mapping of response function gradients. Therefore, the step size should remain in a range where the correction component is small compared to the in-plane component.

2.3 Intrusive coupling of subdomains

Although the semi-intrusive method presented in Sect. 3 does not rely on intrusive coupling, the Lagrange multiplier approach serves as a reference for validation and is therefore briefly described here.

Two (or more) non-conforming meshes can be coupled through the Lagrange multiplier method (Barlow 1982). The augmented stiffness matrix, $\tilde{\mathbf{K}}$, is composed of the stiffness contribution originating from the base, $\mathbf{K}_b$, the stiffeners, $\mathbf{K}_s$, and a set of linear coupling relations encapsulated in matrix, $\mathbf{G}$.

$$\tilde{\mathbf{K}} \tilde{\mathbf{u}} = \tilde{\mathbf{f}} \quad (6a)$$

$$(\mathbf{K}_b \oplus \mathbf{K}_s(\mathbf{x}_i) \oplus \mathbf{G}(\mathbf{x}_i)) \tilde{\mathbf{u}} = \tilde{\mathbf{f}}. \quad (6b)$$

The stiffener node coordinates $\mathbf{x}_i$ serve as design variables, leaving $\mathbf{K}_b$ independent of them. Since base and stiffeners do not share nodes, $\tilde{\mathbf{K}}$ is predominantly diagonal (cf. Eq. (7)). For n_s stiffeners, $\mathbf{K}_s$ decomposes into individual

components $\mathbf{K}_s^{n_s}$ along the diagonal, with only $\mathbf{G}$ providing coupling between domains.

$$\tilde{\mathbf{K}} = \begin{bmatrix} \mathbf{K}_b & 0 & & & \\ & \begin{bmatrix} \mathbf{K}_s^1 & 0 \\ & \ddots \\ 0 & \mathbf{K}_s^{n_s} \end{bmatrix} & & & \\ & & \mathbf{G} & & \\ & & & 0 & \end{bmatrix}. \quad (7)$$

The coupling between the stiffener nodes and the base nodes is established using linear equations that relate the displacements of these nodes (Zienkiewicz et al. 2013).

To ensure compatibility of displacements, the dependent node of the stiffener is connected to the corresponding base element nodal displacements through isoparametric weight coefficients. This can be expressed as:

$$\mathbf{g}^i = \sum_{j=1} N_j \mathbf{u}_b^j - \mathbf{u}_s^i = 0, \quad (8)$$

where $\mathbf{u}_s^i$ is the i -th stiffener node displacement, $\mathbf{u}_b^j$ is the j -th base element nodal displacement, and N_j are the shape function values of the base element evaluated at the stiffener node coordinates. These coupling equations then assemble into matrix $\mathbf{G}$ as in Eq. (9):

$$\begin{bmatrix} \mathbf{g}^1(\mathbf{x}_i^1) \\ \vdots \\ \mathbf{g}^{n_s}(\mathbf{x}_i^{n_s}) \end{bmatrix} = \mathbf{G}(\mathbf{x}_i) \begin{bmatrix} \mathbf{u}_b \\ \mathbf{u}_s \end{bmatrix} = 0. \quad (9)$$

This formulation requires explicit knowledge of the element shape functions, a key limitation addressed by the semi-intrusive approach in Sect. 3.

3 Semi-intrusive approach for stiffener path optimisation

This section presents the proposed non-intrusive method for coupling non-conforming stiffener and base meshes within a commercial FEA workflow. The core idea is to replace the Lagrange multiplier coupling, which requires explicit shape function access, with a local element subdivision scheme that achieves equivalent coupling through shared nodal connectivity. The scheme is first described and validated against the intrusive reference, then embedded into the complete shape optimisation loop.




Fig. 3 Element subdivision on 1 node

3.1 Element subdivision scheme

As established in Sect. 2, Lagrange multiplier coupling requires explicit knowledge of element shape functions, unavailable in commercial software and non-differentiable for gradient-based optimisation. Our approach exploits a fundamental FEM property: nodal connectivity implies displacement coupling. By locally subdividing the base element into triangular elements that share the stiffener bottom node(s) as common node(s), we enforce coupling without accessing shape function formulations. The stiffener node may lie at any position in the interior of a base element, and the subdivision is formed by connecting that node to the surrounding corner nodes (cf. Fig. 3). Stiffener nodes are processed one at a time, and the list of searchable base elements is rebuilt before each node is located. Because subdividing an element removes the original and replaces it with new triangles, two stiffener nodes initially sharing a base element are still handled correctly: the first subdivides it, and the second is then found in the triangle it now falls into rather than in the element that no longer exists, so the search always operates on the current mesh.

To ensure that the element subdivision scheme does not significantly affect the quality of the response function, we compare the compliance of a square domain (40×40 mm) under three different load cases: (1) Half MBB beam, (2) Simply supported plate with central point load, and (3) Cantilever plate with torsional loading. The loads (red arrows) represent forces of $F = 100$ N applied as illustrated in Table 1. These benchmarks consider the base geometry alone, without stiffeners, and all elements, both the quadrilaterals and the triangles, are linear.

The material properties are Young's modulus $E = 210$ GPa, Poisson's ratio $\nu = 0.3$, and a uniform thickness of 1.5 mm. The unsplit case consists of 100 regular square elements arranged in a 10×10 grid, providing a uniform discretisation of the domain.

The split case, by contrast, subdivides base elements at random nodal positions representative of a possible stiffener layout, introducing irregular triangular elements into the otherwise structured mesh. A representative instance of such an arbitrarily subdivided mesh is shown in Fig. 4, with the




Fig. 4 Split mesh configuration with elements subdivided at random nodes representing a possible stiffener layout. Detail views show a two-node subdivision case and elements with poor aspect ratios

Table 1 Compliance comparison between unsplit and split mesh configurations for the three benchmark load cases

Load case			
Compliance			
- Unsplit mesh	0.242354	4.863126	6.872827
- Split mesh	0.241138	4.868793	6.870582
Relative error	0.504%	0.116%	0.033%

detail views highlighting two particularly challenging configurations, namely a two-node subdivision and elements with poor aspect ratios. Even these worst-case geometries introduce negligible deviation from the unsplit mesh response.

The compliance of the system is evaluated for both the unsplit (as it would happen if using the Lagrange multiplier coupling) and split (element subdivision) configurations, and the results are compared in Table 1.

The average relative error across all three load cases is 0.22%, which is well within the acceptable tolerance for engineering applications. This negligible discrepancy demonstrates that the element subdivision scheme preserves the structural response characteristics of the system with high fidelity. Consequently, the proposed non-intrusive coupling approach is validated for use in structural optimisation, providing accurate results whilst maintaining compatibility with commercial finite element software.

The subdivision serves solely for solving the structural primal with the commercial FE software (cf. Fig. 5) and the shape updates are always applied to the original unsplit mesh. In Appendix A, other subcases and mesh-quality safeguards that were implemented, including the mesh-quality controls that eliminate or merge degenerate elements, are also showcased. These controls keep the compliance response accurate,

but the small, irregular triangular sub-elements they cannot fully remove can still introduce localised stress spikes. This is illustrated for the non-uniform geometry benchmark in Appendix B. Beyond this local effect, the overall base and stiffener mesh resolution could also influence the optimised shape and its compliance. Appendix C examines this directly through a mesh sensitivity analysis on the two-stiffener cantilever benchmark introduced later in Sect. 5.

3.2 Implementation for shape optimisation

The proposed semi-intrusive framework for stiffener path optimisation, built around the non-intrusive coupling approach, is illustrated in Fig. 5. First, the base elements are subdivided at stiffener nodal positions to establish coupling. Structural primal and adjoint analysis is performed using the commercial FEA software Altair OptiStruct, followed by the backward and forward filtering of the vertex morphing parametrisation (cf. Sect. 2.2) with the shape update being carried onto the unsplit mesh file. This unsplit mesh is maintained as a copy in parallel with the subdivided one and receives the updated coordinates only at the mapping stage, so each new iteration starts from a clean base mesh whose nodal connectivity is rebuilt by locating the stiffener nodes and subdividing the base elements again. This cycle repeats until convergence criteria are met, preserving the non-conformity of the meshes and allowing unrestricted stiffener node movement throughout the optimisation.

4 Manufacturing constraints

4.1 Stiffener parametrisation via layered normal-based discretisation

Whilst 1D beam elements could represent stiffeners, defining cross-section orientations in commercial FE software becomes problematic for non-planar geometries typical of automotive components. Therefore, the stiffener mesh is constructed using 2D shell elements with a hierarchical two-level node structure. Anchor nodes form the bottom row in contact with the base mesh and serve as design variables. Link nodes are positioned along the local surface normal at prescribed offset distances, forming the stiffener height (cf. Fig. 6). This ensures stiffeners remain orthogonal to the base geometry throughout optimisation, a critical requirement for WAAM-manufactured reinforcements.

When anchor nodes receive shape updates $\delta \tilde{\mathbf{x}}_i$, link nodes must follow whilst preserving orthogonality towards the base. The predicted link position is computed by projecting

Fig. 5 Semi-intrusive stiffener path optimisation framework. The base mesh is subdivided at stiffener nodal positions to establish coupling, structural analysis is performed on the subdivided mesh, and the resulting shape update is applied back to the original unsplit mesh. This repeats until convergence





Fig. 6 Layered normal-based stiffener discretisation: anchor nodes $\mathbf{x}_i$ define the stiffener path, whilst link nodes $\mathbf{x}_\ell$ are positioned along the surface normal at height h

along the updated normal direction:

$$\mathbf{x}'_\ell = \mathbf{x}'_i + h \frac{\mathbf{n}'_i}{\|\mathbf{n}'_i\|} \quad (10)$$

where $\mathbf{x}'_i$ is the updated anchor position after in-plane projection, h is the distance between anchor and link nodes, and $\mathbf{n}'_i$ is the surface normal at the new anchor location, interpolated from the base surface nodal normals rather than taken from the single element on which the anchor lands, which keeps the direction accurate on curved geometry. The link node update is then:

$$\Delta \mathbf{x}_\ell = \mathbf{x}'_\ell - \mathbf{x}_\ell \quad (11)$$

This update is fully determined by the updated anchor position, and $\Delta \mathbf{x}_\ell$ is the update then applied to the link node. This parametrisation reduces the design space to anchor nodes only, implicitly preserves stiffener height, and naturally follows surface curvature throughout optimisation. Similarly to the in-plane projection, this parametrisation introduces a non-linear relation in the shape update of anchor and link nodes that cannot be accounted for in the mapping of response function gradients because the sensitivities evaluated at the link nodes are projected back onto the anchor nodes. Consequently, an overly large step size can introduce zigzagging into the optimisation. Again, a relatively small step size should be used, in a range where the anchor and link node updates do not differ extensively.

4.2 Angle-based penalised minimum distance constraint

To ensure WAAM manufacturability, a penalised minimum distance constraint enforces distance between disconnected stiffener components. When stiffeners are printed too close, thermal interaction during manufacturing combined with increased local thermal input may create major thermal deflection, rendering the part non-manufacturable or dimensionally non-compliant. The constraint is made angle-aware so that crossings are permitted, since applying it unconditionally would excessively restrict the design space without a corresponding manufacturability benefit.

4.2.1 Formulation

The reference surface, i.e. the base geometry, contains multiple disconnected stiffeners. Computing distances directly between discrete node pairs would not reliably capture the true minimum distance, since the closest point on a neighbouring stiffener may lie between nodes rather than at a node itself. To address this, we employ a point-to-element distance function, where each node x_i is measured against the continuous line elements that describe the neighbouring stiffener curves. The unsigned distance to the closest point y^* on the nearest element e^* is:

$$d_i = \min_{y^* \in e^*} \|x_i - y^*\|_2, \quad (12)$$

where y^* is obtained by orthogonal projection onto the element's parametric surface. A violation exists when $d_i < d_{\min}$, and its magnitude is:

$$\Delta_i = d_{\min} - d_i > 0. \quad (13)$$

To preserve design freedom whilst managing thermal distortion, the constraint distinguishes between brief crossings and extended parallel proximity. Stiffener crossings involve only localised thermal interaction and are permitted to enable topological flexibility during optimisation. However, extensive lengths where stiffeners run parallel at close proximity create prolonged thermal interaction that must be prevented. An angle-based penalisation factor $w(\theta)$ modulates the constraint based on surface normal parallelism. The parallelism between the unit normals $\hat{n}_i$ and $\hat{n}_j$ at the two interacting surfaces is measured via the absolute cosine,

$$|\cos \theta| = |\hat{n}_i \cdot \hat{n}_j|, \quad \theta \in [0^\circ, 90^\circ], \quad (14)$$

using the absolute value to handle arbitrary mesh normal orientations on disconnected components, where $\theta = 0^\circ$ indicates perfectly parallel surfaces and $\theta = 90^\circ$ indicates perpendicular surfaces.

A natural first candidate for the weight is $w = |\cos \theta|$ clamped to zero above a threshold θ_t . However, since $\cos \theta_t \neq 0$ in general, this produces a finite jump discontinuity at the activation boundary:

$$\lim_{\theta \rightarrow \theta_t^-} |\cos \theta| = \cos \theta_t \neq 0. \quad (15)$$

Whenever a node pair's angle crosses the threshold between successive optimisation iterations, the response value jumps




Fig. 7 Point-to-element distance and angle computation: x_i is the evaluated node, y^* the closest point on the nearest element e^* , and θ the angle between unit surface normals $\hat{n}_i$ and $\hat{n}_j$

discontinuously, which can stall or destabilise gradient-based convergence. A well-posed weight must instead satisfy the C^1 conditions ensuring both the value and slope vanish continuously at the threshold. The raised cosine (Hann window) satisfies these requirements:

$$w(\theta) = \begin{cases} \frac{1}{2} \left(1 + \cos \left(\frac{\pi \theta}{\theta_t} \right) \right) & \theta < \theta_t, \\ 0 & \theta \geq \theta_t, \end{cases} \quad (16)$$

achieving $w(0) = 1$ for perfectly parallel configurations and decaying smoothly to zero at the threshold with zero slope ($w'(\theta_t) = 0$), ensuring C^1 continuity across the activation boundary. A visual representation of the distance and angle computation is provided in Fig. 7.

4.2.2 Aggregation and sensitivity analysis

For each node x_i , the weighted violation combines the distance violation with the penalisation weight:

$$v_i = \Delta_i \cdot w(\theta_i) = (d_{\min} - d_i) w(\theta_i). \quad (17)$$

These are aggregated using a p -norm ($p = 4$) to provide a smooth, differentiable constraint:

$$g = \left(\sum_{i \in \mathcal{V}} v_i^p \right)^{1/p}, \quad (18)$$

where $\mathcal{V} = \{i : d_i < d_{\min}\}$ denotes the active violation set. The p -norm provides a smooth, differentiable approximation to the maximum violation, remaining everywhere differentiable and avoiding the sub-gradient issues of the exact maximum.

Sensitivities are derived analytically using the frozen-normals approximation: surface normals are treated as

constant parameters within each gradient evaluation and recomputed from the updated geometry at the start of each iteration, consistent with a standard semi-analytical (lagged-Jacobian) scheme. The gradient of g with respect to a node position $\mathbf{x}$ follows from the chain rule:

$$\frac{\partial g}{\partial \mathbf{x}} = \frac{1}{g^{p-1}} \sum_{i \in \mathcal{V}} v_i^{p-1} \frac{\partial v_i}{\partial \mathbf{x}}. \quad (19)$$

Expanding $v_i = \Delta_i \cdot w_i$ via the product rule gives two contributions:

$$\frac{\partial v_i}{\partial \mathbf{x}} = w_i \frac{\partial \Delta_i}{\partial \mathbf{x}} + \Delta_i \frac{\partial w_i}{\partial \mathbf{x}}. \quad (20)$$

The second term requires differentiating the weight w_i , which depends on the surface normals $\hat{\mathbf{n}}_i$ and $\hat{\mathbf{n}}_j$, with respect to the node coordinates:

$$\frac{\partial w_i}{\partial \mathbf{x}} = \frac{\partial w_i}{\partial \hat{\mathbf{n}}_i} \frac{\partial \hat{\mathbf{n}}_i}{\partial \mathbf{x}} + \frac{\partial w_i}{\partial \hat{\mathbf{n}}_j} \frac{\partial \hat{\mathbf{n}}_j}{\partial \mathbf{x}}. \quad (21)$$

Since surface normals are computed as area-weighted averages of surrounding element normals, evaluating $\partial \hat{\mathbf{n}} / \partial \mathbf{x}$ exactly requires differentiating the area-weighted cross-product through the full mesh connectivity, a dedicated mesh sensitivity implementation that is not available in the present framework. The frozen-normals approximation therefore sets

$$\frac{\partial w_i}{\partial \mathbf{x}} \approx 0, \quad (22)$$

treating $\hat{\mathbf{n}}$ as a constant parameter within each gradient evaluation. Normals are recomputed from the updated geometry at the start of each optimisation iteration, so the approximation is consistent with a standard semi-analytical (lagged-Jacobian) scheme. Equation (20) then reduces to:

$$\frac{\partial v_i}{\partial \mathbf{x}} \approx w_i \frac{\partial \Delta_i}{\partial \mathbf{x}}. \quad (23)$$

The neglected second term is suppressed by three complementary mechanisms. First, the optimisation drives $\Delta_i \rightarrow 0$, so the term vanishes wherever the constraint is satisfied. Furthermore, the Hann window satisfies $w'(\theta_i) = 0$, so $\partial w_i / \partial \mathbf{x}_k$ vanishes at the activation boundary, where the error would otherwise be largest. Lastly, for intermediate angles where neither factor is small, the lagged-Jacobian update recomputes normals each iteration, so any residual gradient error

at the current iteration is corrected when w is re-evaluated on the next one.

For the point-to-element distance, treating the closest point $\mathbf{y}^*$ as fixed during differentiation yields:

$$\frac{\partial d_i}{\partial \mathbf{x}_i} = \frac{\mathbf{x}_i - \mathbf{y}^*}{d_i} \Rightarrow \frac{\partial \Delta_i}{\partial \mathbf{x}_i} = -\frac{\mathbf{x}_i - \mathbf{y}^*}{d_i}. \quad (24)$$

Combining all terms, the gradient of the p -norm constraint is:

$$\frac{\partial g}{\partial \mathbf{x}} \approx \frac{1}{g^{p-1}} \sum_{i \in \mathcal{V}} v_i^{p-1} w_i \frac{\partial \Delta_i}{\partial \mathbf{x}}, \quad (25)$$

which for $\partial \Delta_i / \partial \mathbf{x}_i = -(\mathbf{x}_i - \mathbf{y}^*) / d_i$ yields:

$$\frac{\partial g}{\partial \mathbf{x}_i} = \frac{1}{g^{p-1}} \sum_{i \in \mathcal{V}} v_i^{p-1} w_i \frac{\mathbf{y}^* - \mathbf{x}_i}{d_i}, \quad (26)$$

weighting each violation by its relative magnitude.

The manufacturing constraints introduced in this section and in Sect. 2, namely stiffener orthogonality, minimum distance, packaging and mass, act purely at the geometric level. Their benefit is demonstrated in Sect. 5 by comparing optimised layouts obtained with and without each constraint active. Thermal effects intrinsic to WAAM, such as build-up distortion and residual stresses, are not modelled. Even in topology optimisation, where the design representation is comparatively simple, embedding such effects has been shown to require dedicated build-process simulations and adjoint sensitivities, and, even after simplifying the underlying thermal history through an inherent-strain approximation, to increase computation time by roughly two orders of magnitude whilst degrading convergence (Misiun et al. 2021). Given this cost, and since a comparably accurate experimental validation would require the same finer modelling, we regard thermal distortion and residual-stress constraints as an important direction for future work. The present, purely geometric constraints, by contrast, are verified directly at the end of every optimisation run and therefore need no separate experimental confirmation that the manufacturing team's requirements are met.

5 Numerical results

This section evaluates the proposed framework through four benchmark problems of increasing complexity: a flat cantilever plate, a curved cylindrical shell section, a plate with a central circular cut-out, and a non-uniform irregular geometry. Each case targets a specific aspect of the




Fig. 8 Geometry of 2D cantilever plate

method: coupling accuracy, curved-surface handling, geometric adaptability and full-constraint integration. They are also compared against reference results from the literature. The objective convergence histories for all cases are collected in Appendix D, and the corresponding constraint convergence histories in Appendix E.

5.1 Cantilever

To validate the proposed non-intrusive coupling approach, we consider a classical 2D cantilever plate problem based on the benchmark presented by Bojczuk and Szeleblak (2008). The rectangular plate has dimensions 600 mm $\times$ 400 mm and is fully clamped along its left edge. A concentrated load $F = 5$ kN is applied at the centre of the right edge. The base plate and stiffeners are modelled with shell elements, where the stiffeners have a thickness of 10 mm and a height of 20 mm. The material properties correspond to an aluminium alloy with Young's modulus $E = 7.5 \times 10^4$ MPa and Poisson's ratio $\nu = 0.35$. The problem setup is illustrated in Fig. 8.

Following the reference paper, we tested two configurations: one with two stiffeners and another with four stiffeners. The four-stiffener case was additionally evaluated both with and without the distance constraint to assess the impact of enforcing minimum distance between stiffeners.

5.1.1 Two stiffeners

The optimisation successfully converged to a final stiffener configuration closely resembling the results reported in Bojczuk and Szeleblak (2008), as shown in Fig. 9. Despite starting from a significantly different initial guess, the optimiser was able to navigate the design space and reach an equivalent optimal shape, demonstrating the robustness of the proposed method.




Fig. 9 Cantilever plate with two stiffeners: initial—grey, and optimised stiffener shapes—blue




Fig. 10 Cantilever plate with four stiffeners: initial—grey, and optimised stiffener shapes—blue

5.1.2 Four stiffeners

The optimised configuration again achieves a shape consistent with the results presented in Bojczuk and Szeleblak (2008). Notably, in the reference paper, the authors locked the two outer stiffeners (obtained from their previous two-stiffener optimisation) and only optimised the two internal stiffeners. In contrast, our approach performed a simultaneous optimisation of all four stiffeners without constraining any of them, yet still converged to a comparable final shape. This demonstrates the method's ability to explore a larger design space whilst arriving at solutions of equivalent quality, as shown in Fig. 10.

This case also serves as an ideal benchmark for validating the angle-based penalised minimum distance constraint (Sect. 4.2), configured with a minimum distance of 40 mm and a threshold angle of 45°.

The results in Fig. 11 clearly show the distance constraint being triggered for near-parallel stiffener segments (highlighted in red), enforcing the minimum distance, whilst crossing sections (marked in green) remain unaffected, as desired.




Fig. 11 Cantilever plate with four stiffeners (with distance constraint) optimised shape




Fig. 12 Geometry of cylindrical shell section (force is applied in the centre point)

5.2 Cylindrical section

To demonstrate the framework's capability on curved geometries, we consider a cylindrical shell section benchmark following the problem setup of Feng et al. (2021). The shell has a radius $R = 20$ mm, subtends an angle of $\theta = 2$ rad, and has a length $L = 40$ mm. The four corners are simply supported, and a concentrated load $F = 1$ kN is applied at the centre of the shell, pointing downward. The problem setup is illustrated in Fig. 12.

This benchmark enables comparison with existing literature results and validates the in-plane projection mechanism on curved surfaces.

The optimisation on the curved cylindrical geometry, shown in Fig. 13, yields results comparable to those reported by Feng et al. (2021). Despite the reference study employing a topology optimisation-based method, both approaches converge to a similar characteristic shape, a cross pattern centred at the applied load with branches extending towards each of the four corner supports. This demonstrates the framework's capability to handle non-planar surfaces and validates the in-




Fig. 13 Cylindrical shell section: initial—grey, and optimised stiffener shapes—blue




Fig. 14 Geometry of rectangular plate with central circular hole under point load at centroid, simply supported at each corner

plane projection mechanism for curved shells. Furthermore, the optimised solution exhibits stiffener orthogonality, confirming the effectiveness of the stiffener parametrisation via layered normal-based discretisation.

5.3 Plate with centre cut-out under point load at centroid

To evaluate the framework's performance on geometries with internal features, we consider a rectangular plate with a central circular hole subjected to a point load at its centroid. The plate has dimensions 40 mm $\times$ 60 mm and contains a circular hole of diameter $d = 10$ mm at its centre. The plate is simply supported at each of its four corners, whilst a concentrated force $F = 500$ N is applied at the centroid. The base plate has a thickness of 1.5 mm and the stiffeners have a thickness of 4 mm. The material is stainless steel with Young's modulus $E = 210$ GPa and Poisson's ratio $\nu = 0.3$. The problem setup is illustrated in Fig. 14.

This benchmark problem introduces a challenging test case for stiffener path optimisation where stiffeners must navigate around the internal feature.

For this loading configuration with corner supports and a central point load, the expected optimal stiffener layout shown in Fig. 15 would be a cross pattern centred at the load and extending towards each corner, similar to the cylindrical shell case. However, the presence of the central cut-out prevents stiffeners from passing through that region, both physically (no material exists there) and methodologically (the subdivision scheme requires a base element to subdivide). Consequently, the optimiser produces a modified layout that routes around the hole whilst still efficiently trans-




Fig. 15 Plate with central circular hole: initial—grey, and optimised stiffener shapes—blue




Fig. 16 Non-uniform geometry

ferring loads to the corner supports. This confirms that the method adequately adapts to geometric features and provides optimised shapes that account for them.

5.4 Non-uniform geometry

This benchmark demonstrates the optimiser's capability to handle non-uniform and irregular geometries without difficulties. Since real-world applications typically involve complex part shapes rather than idealised rectangular domains, this feature is of extreme relevance for practical engineering applications. Beyond geometry handling, this case also serves as a full integration test in which all constraints and parametrisations, the layered normal-based parametrisation, the in-plane projection, and the angle-based penalised distance constraint, are active simultaneously, allowing us to verify that their interactions produce consistent and stable optimisation behaviour. The base plate has a thickness of 1.5 mm and the stiffeners have a thickness of 4 mm. The material is stainless steel with Young's modulus $E = 210$ GPa and Poisson's ratio $\nu = 0.3$. The geometry is shown in Fig. 16.

The results on the non-uniform geometry in Fig. 17 demonstrate that the method can effectively handle irregular




Fig. 17 Non-uniform geometry: initial—grey, and optimised stiffener shapes—blue

shapes. The stiffeners adequately move along the reference geometry whilst preserving orthogonality, confirming that the layered normal-based discretisation remains valid for complex domains. This benchmark validates the framework's applicability to realistic industrial components where idealised rectangular geometries are rarely encountered. It is worth noting that cases involving large deviations from the initial layout, such as complex non-uniform geometries, may, under specific conditions, exhibit bridging artefacts when neighbouring nodes drift beyond the filter radius, losing their smoothness coupling, as discussed in Appendix F. This is not a fundamental limitation of the method, but rather an aspect to be aware of when the initial and optimised designs differ significantly.

6 Conclusions

The presented semi-intrusive shape optimisation framework for stiffened structures demonstrates several key capabilities:

- The element subdivision scheme provides a non-intrusive coupling between non-conforming stiffener and base meshes without requiring access to proprietary shape function formulations, making the method compatible with commercial FEA software and existing simulation workflows. The average compliance error between the split mesh and the unsplit mesh remained below 0.6% across all tested load cases, confirming that the

subdivision introduces negligible error into the structural response. A dedicated mesh sensitivity analysis (Appendix C) further showed that refining the base and stiffener mesh from 8 mm down to 1 mm changes the optimised compliance by only about 2 %, with the 4 mm mesh used throughout Sect. 5 already within 1 % of the finest mesh tested and the resulting stiffener layouts remaining consistent across all resolutions, confirming that the method is not overly sensitive to mesh size.

- The cantilever and cylindrical shell benchmarks achieved final configurations comparable to those reported in the reference literature, confirming that the non-intrusive coupling not only preserves structural accuracy but also remains fully compatible with gradient-based shape optimisation, allowing the method to reach the expected optima.
- The angle-based penalised minimum distance constraint successfully enforced minimum distance between near-parallel stiffener segments whilst leaving crossing sections unaffected. This behaviour was validated on the four-stiffener cantilever case and demonstrates the constraint's ability to address WAAM thermal interaction requirements without unnecessarily restricting path flexibility during optimisation. Across all benchmarks, the packaging constraint also approached its bound monotonically, without the violate–correct oscillation typical of the plain RGP method (Appendix E), confirming the smoother active constraint handling attributed to the QNRGP algorithm in Sect. 2.
- The layered normal-based parametrisation maintained stiffener orthogonality to the base surface throughout all optimisations, confirming its validity as a WAAM-oriented parametrisation. The in-plane projection correctly constrained stiffener paths to follow the reference surface, as demonstrated on the cylindrical shell benchmark where results matched those reported by topology optimisation-based methods.
- The framework successfully handles geometries with internal features (plate with central hole) and irregular non-uniform shapes. The non-uniform case additionally served as a full integration test, with all constraints and parametrisations active simultaneously, confirming consistent and stable interactions between components. Together, these cases validate the framework's applicability to realistic industrial components where idealised rectangular geometries are rarely encountered.
- The proposed manufacturing constraints act purely at the geometric level and are verified directly at the end of every optimisation run, requiring no separate experimental validation. Thermal effects intrinsic to WAAM, such as build-up distortion and residual stresses, are not modelled and are left as an important direction for future work (Sect. 4).

Acknowledgements The authors wish to thank the BMW ShapeModule project and especially Mr. Muneeb Rehman (BMW Group, Munich, Germany) for the technical support.

Author Contributions Duarte F. Valoroso Madeira developed the methodology, produced results, and wrote the manuscript. Armin Geiser developed the methodology and reviewed the manuscript. Majid Hojjat supervised and reviewed the manuscript. Benedikt Kriegesmann supervised and wrote the manuscript.

Funding Open Access funding enabled and organized by Projekt DEAL.

Data Availability No data have been produced within this research.

Declarations

Conflict of interest The authors declare that they have no conflict of interest.

Replication of results The proposed methods are implemented in the optimisation framework BMW ShapeModule (shapemodule@bmw.de) (no public access).

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

Appendix A Subdivision schemes and numerical safeguards

For simplification purposes, the base mesh element will always be treated as a first-order quadrilateral element, but the method should work for other shell elements (with the necessary modifications). Because a stiffener node can fall at an arbitrary interior position of a base element, away from its vertices and edges, the same subdivision applies regardless of where the node lands. For the complex, curved geometries found in industrial practice, biquadratic elements are a legitimate alternative, since their curved edges follow a curved reference surface more closely and a sufficiently fine subdivision would in principle capture it. A direct subdivision, however, would embed linear sub-elements within a quadratic surrounding mesh, which is inconsistent. Subdividing into quadratic sub-elements instead is possible, but it enlarges the stiffness matrix: splitting a single quadratic triangle into six quadratic sub-triangles already adds six nodes




Fig. 18 Element subdivision on 2 nodes

to the base mesh, each carrying six degrees of freedom, and hence 36 rows and columns in the stiffness matrix, which becomes prohibitively expensive for large industrial models. Given this higher dimensionality and computational cost, first-order elements are preferred here. A NURBS-based discretisation could likewise be argued for, but as the method is intended to remain semi-intrusive and usable with commercial software, the retained option must avoid any need to access the element shape functions.

Given that two stiffener nodes frequently fall within the same base element, a subdivision scheme for this case is also considered (cf. Fig. 18).

This second subdivision scheme simply adds an extra subdivision on one of the created triangles, effectively subdividing as if for only one node. On the next node, it again subdivides the base element below it, but in this case only into three sub-elements because it operates on a triangle. This approach effectively allows for an unlimited number of nodes within a single initial quadrilateral base element. However, this is generally undesirable because it would produce very small elements that exhibit numerical issues. Although it is theoretically possible for many stiffener nodes to fall within the same base element, this situation must be avoided in practice by enforcing a minimum distance constraint between stiffener nodes that are topological neighbours. Regardless, we assume this as a known limitation of the current implementation.

To avoid numerical issues due to ill-conditioned elements (i.e. very distorted elements), which lead to zero or negative Jacobians, two additional safeguards are employed. In cases where the subdivision produces an element with area smaller than 0.1% of the original base element, that element is eliminated to prevent numerical instabilities whilst avoiding a significant topological change (cf. Fig. 19). Alternatively, when two subdivided elements have such small areas, which implies the stiffener node is very close to a corner of the base element, the corner node is repositioned to the stiffener node location, thereby preserving the mesh topology (cf. Fig. 20). These same safeguards are also applied in the two-node subdivision case.




Fig. 19 Element elimination: when a subdivided element has area smaller than 0.1% of the original, it is removed (shown in black)




Fig. 20 Node repositioning: when the stiffener node is close to a corner and two elements become too small, the corner node is moved to the stiffener position, preserving topology

Appendix B Local stress concentrations

As discussed in Sect. 3.1, the element subdivision scheme can produce small, distorted triangular sub-elements, and, being constant-strain elements, these resolve stress less accurately than the surrounding quadrilaterals. This is fundamentally a numerical artefact of the discretisation rather than a physical effect. Most of the observed spikes occur at the stiffener ends, plausibly because the abrupt change in bending stiffness there locally concentrates curvature, making this numerical resolution loss more visible at these locations. Figure 21 illustrates this on the non-uniform geometry benchmark of Sect. 5, where the von Mises stress field exhibits a localised spike at a stiffener end. This does not affect the stability of the optimisation, whose convergence for this benchmark remains smooth throughout, as shown in Appendix D (cf. Fig. 26). Since the objective and constraints used throughout this work are compliance- and geometry-based rather than stress-based, this is a known but acceptable limitation, as a stress-based optimisation formulation was always outside the scope of this work.

Appendix C Mesh sensitivity analysis

Since the coupling is established locally through the element subdivision scheme of Sect. 3.1, the overall base and stiffener mesh resolution could, in principle, affect both the optimised shape and its compliance. To assess this, the two-stiffener




Fig. 21 Von Mises stress field for the non-uniform geometry benchmark, showing a localised stress spike at the end of a stiffener

Table 2 Compliance of the two-stiffener cantilever benchmark for four uniform mesh sizes, with the vertex morphing filter radius kept constant

Mesh size [mm]	8	4	2	1
Compliance	329.3	324.9	321.4	321.8

cantilever benchmark of Sect. 5 was re-optimised using four uniform mesh sizes, 8, 4 (the size used for the results reported in Sect. 5), 2 and 1 mm, applied consistently to both the base and the stiffeners, whilst the vertex morphing filter radius was kept fixed across all four runs to isolate the effect of the mesh from that of the filter. The resulting optimised shapes, compared in Fig. 22, are visually very similar across all four mesh sizes, with only the coarsest, 8 mm, mesh showing a slightly more noticeable deviation, as expected, indicating that the optimiser converges to essentially the same stiffener layout regardless of resolution. The final compliance values, summarised in Table 2, decrease with mesh refinement down to the 2 mm mesh, with the relative change between successive mesh sizes shrinking from 1.34 % (8 to 4 mm) to 1.08 % (4 to 2 mm), followed by a marginal 0.12 % increase between the 2 mm and 1 mm meshes. This overall trend is the expected asymptotic convergence behaviour of a standard finite element discretisation under mesh refinement rather than an artefact of the element subdivision scheme, and the small non-monotonic step between the two finest meshes is consistent with the optimiser settling into a slightly different, comparably optimal local minimum rather than with mesh-dependent degradation. Overall, this indicates that the method is not overly sensitive to mesh size. In particular, the compliance obtained with the 4 mm mesh used throughout Sect. 5 lies within 1 % of the value obtained with the finest, 1 mm, mesh, supporting the mesh resolution adopted for the reported results.




Fig. 22 Optimised stiffener shapes for the two-stiffener cantilever benchmark at four mesh sizes: 8 mm (green), 4 mm (blue), 2 mm (yellow) and 1 mm (red)




Fig. 23 Convergence history of the relative compliance change for the cantilever benchmark with 2 stiffeners

Appendix D Objective convergence histories

The following figures (cf. Figs. 23, 24, 25, 26 and 27) show the relative change of the objective function with respect to the first iteration for each optimisation case. In all examples, a clear reduction of the objective relative to the initial design is observed throughout the optimisation process. The optimisation algorithm verifies convergence by examining the Karush–Kuhn–Tucker (KKT) conditions, that is, by monitoring the norm of the objective gradient projected onto the tangent subspace of the currently active constraints. Since the problem involves only inequality constraints, written generically as $g_j(\mathbf{x}) \leq 0$ with associated Lagrange multipliers λ_j (here the packaging, minimum distance and mass responses of Eq. (1)), complementary slackness ($\lambda_j g_j(\mathbf{x}) = 0$) removes the multipliers of all inactive constraints, so a vanishing projected gradient is equivalent to the reduced stationarity condition,

$$\nabla f(\mathbf{x}) + \sum_{j \in \mathcal{A}} \lambda_j \nabla g_j(\mathbf{x}) = 0, \quad (27)$$

on the active set $\mathcal{A} = \{j : g_j(\mathbf{x}) = 0\}$, leaving no feasible improving direction. The tolerance for subsequent iterations is calculated as a user-defined percentage of the norm of this projected direction in the initial iteration. This percentage is set to 5% in all presented examples.

On the 4 stiffener configuration we have the only case where two directly comparable optimised configurations




Fig. 24 Convergence history of the relative compliance change for the cylinder benchmark




Fig. 25 Convergence history of the relative compliance change for the plate with hole benchmark




Fig. 26 Convergence history of the relative compliance change for the non-uniform load benchmark




Fig. 27 Convergence history comparison of the relative compliance change for the cantilever benchmark with 4 stiffeners, with and without minimum distance constraint

exist (with and without the distance constraint), which offers a unique opportunity to quantify the trade-off between structural performance and manufacturability. As expected, the constrained solution yields a slightly higher compliance than its unconstrained counterpart (cf. Fig. 27), however, the difference remains below 5%.

In contrast, the gain in manufacturability is substantially larger. Whilst it cannot be expressed as a single numerical metric, the unconstrained layout contains near-parallel stiffener segments that violate the minimum spacing




Fig. 28 Packaging constraint history for the cantilever benchmark with two stiffeners




Fig. 29 Packaging constraint history for the cantilever benchmark with four stiffeners, with and without the minimum distance constraint

required for WAAM deposition, rendering it effectively non-manufacturable, whereas the constrained layout fully respects the geometric requirements and can be directly fabricated. The improvement in manufacturability therefore far exceeds the modest loss in structural efficiency, making the constrained solution the clearly preferable choice in practice.

Appendix E Constraint convergence histories

Rather than merely confirming feasibility, these histories give empirical support to two mechanisms introduced earlier. The packaging distances (Figs. 28, 29, 30, 31, and 32) approach the box boundary monotonically and settle onto it without the violate–correct oscillation that the plain RGP method would produce, which supports the smoother active constraint handling attributed to the QNRGP method in Sect. 2. Where the packaging constraint becomes active at convergence, namely the four-stiffener cantilever, the plate with hole and the non-uniform geometry, it is the packaging bound that limits the attainable optimum. For the cylinder, it stays inactive throughout, indicating that its optimum lies in the interior of the admissible region and that the corresponding stiffness gain is not packaging-limited. The plate with hole history additionally begins marginally infeasible and is restored to the feasible side within a few iterations, showing that the gradient projection recovers feasibility from an infeasible start without destabilising the descent.

The minimum distance violations (Figs. 33 and 34) decay smoothly to zero, which is both consistent with and a prerequisite for the frozen-normals approximation of Sect. 4.2: as




Fig. 30 Packaging constraint history for the cylinder benchmark




Fig. 31 Packaging constraint history for the plate with hole benchmark




Fig. 32 Packaging constraint history for the non-uniform geometry benchmark




Fig. 33 Minimum-distance constraint history for the cantilever benchmark with four stiffeners

$\Delta_i \rightarrow 0$ the neglected weight-sensitivity term vanishes, so the approximation becomes exact precisely where the constraint is enforced, whilst the absence of discontinuous jumps reflects the C^1 continuity built into the angle-based weight.

Appendix F Stiffener “bridging” and element length constraints

In standard vertex morphing, the filtering operation is performed on the current mesh at each iteration, coupling nodes to their neighbours within the filter radius at that instant. As




Fig. 34 Minimum-distance constraint history for the non-uniform geometry benchmark




Fig. 35 Elongated bridging artefacts between stiffeners resulting from neighbouring nodes exceeding the filter radius during large shape updates

the optimisation progresses, however, neighbouring nodes may drift apart beyond the filter radius, at which point they fall outside each other's sphere of influence and the filter no longer enforces smoothness between them, permitting unbounded relative displacement.

Once nodes exceed the filter radius, this loss of smoothness manifests as artificial “bridging” between stiffeners (cf. Fig. 35), where sharp, elongated features connect paths that were initially separate. Such non-physical transitions arise specifically when the optimised design departs significantly from the initial layout—precisely the scenario in which neighbouring nodes are most likely to drift beyond the filter radius.

To mitigate this, a maximum element length constraint can be introduced, preventing excessive elongation of stiffener segments and indirectly preserving the filter coupling between adjacent nodes. The trade-off is that such a constraint limits how far the shape optimisation can deviate from the initial configuration, potentially restricting the achievable design improvement. This limitation is, however, consistent with the intended role of shape optimisation as a fine-tuning tool rather than a mechanism for large-scale redesign. Moreover, the bridging artefacts only arise when the initial and optimised designs are far apart, which itself argues that their occurrence signals a case better addressed through a different initial layout rather than more constrained shape updates. We are currently investigating several use cases where this constraint may be necessary, in order to establish guidelines

for when its implementation is warranted and to quantify its impact on the attainable optimality.

References

- Aage N, Andreassen E, Lazarov BS, Sigmund O (2017) Giga-voxel computational morphogenesis for structural design. *Nature* 550(7674):84–86. <https://doi.org/10.1038/nature23911>
- Akl W, El-Sabbagh A, Baz A (2008) Optimization of the static and dynamic characteristics of plates with isogrid stiffeners. *Finite Elem Anal Des* 44(8):513–523. <https://doi.org/10.1016/j.finel.2008.01.015>
- Amir O, Majdouba A (2024) Mesh-based topology, shape and sizing optimization of ribbed plates. *Struct Multidisc Optim* 67(6):103. <https://doi.org/10.1007/s00158-024-03824-9>
- Antonau I, Hojjat M, Bletzinger K-U (2021) Relaxed gradient projection algorithm for constrained node-based shape optimization. *Struct Multidisc Optim* 63(4):1633–1651. <https://doi.org/10.1007/s00158-020-02821-y>
- Antonau I, Warnakulasuriya S, Bletzinger K-U, Bluhm FM, Hojjat M, Wüchner R (2022) Latest developments in node-based shape optimization using vertex morphing parameterization. *Struct Multidisc Optim* 65(7):198. <https://doi.org/10.1007/s00158-022-03279-w>
- Asl RN, Shayegan S, Geiser A, Hojjat M, Bletzinger K-U (2017) A consistent formulation for imposing packaging constraints in shape optimization using vertex morphing parametrization. *Struct Multidisc Optim* 56(6):1507–1519. <https://doi.org/10.1007/s00158-017-1819-9>
- Barlow J (1982) Constraint relationships in linear and nonlinear finite element analyses. *Int J Numer Methods Eng* 18(4):521–533. <https://doi.org/10.1002/nme.1620180404>
- Bletzinger K-U (2014) A consistent frame for sensitivity filtering and the vertex assigned morphing of optimal shape. *Struct Multidisc Optim* 49(6):873–895. <https://doi.org/10.1007/s00158-013-1031-5>
- Bojczuk D, Szeleblak W (2008) Optimization of layout and shape of stiffeners in 2D structures. *Comput Struct* 86(13–14):1436–1446. <https://doi.org/10.1016/j.compstruc.2007.05.005>
- Borg Costanzi C (2023) Reinforcing and detailing of thin sheet metal using wire arc additive manufacturing as an application in facades. *Mechanik, Werkstoffe und Konstruktion im Bauwesen*, vol 68. Springer, Wiesbaden. <https://doi.org/10.1007/978-3-658-41540-2>
- Carstensen JV, Kim-Tackowiak H, Liang MY (2023) Improving the manufacturability of highly materially restricted topology-optimized designs with mixed integer linear programming. *Eng Struct* 284:115955. <https://doi.org/10.1016/j.engstruct.2023.115955>
- Chu S, Featherston C, Kim HA (2021) Design of stiffened panels for stress and buckling via topology optimization. *Struct Multidisc Optim* 64(5):3123–3146. <https://doi.org/10.1007/s00158-021-03062-3>
- Feng S, Zhang W, Meng L, Xu Z, Chen L (2021) Stiffener layout optimization of shell structures with B-spline parameterization method. *Struct Multidisc Optim* 63(6):2637–2651. <https://doi.org/10.1007/s00158-021-02873-8>
- Hao P, Wang Y, Jin L, Ma S, Wang B (2023) An isogeometric design-analysis-optimization workflow of stiffened thin-walled structures via multilevel NURBS-based free-form deformations (MNFFD). *Comput Methods Appl Mech Eng* 408:115936. <https://doi.org/10.1016/j.cma.2023.115936>
- Hao P, Liu D, Liu H, Feng S, Wang B, Li G (2024) Intelligent optimum design of large-scale gradual-stiffness stiffened panels via multi-level dimension reduction. *Comput Methods Appl Mech Eng* 421:116759. <https://doi.org/10.1016/j.cma.2024.116759>
- Hirschler T, Bouclier R, Duval A, Elguedj T, Morlier J (2019) The embedded isogeometric Kirchhoff-Love shell: from design to shape optimization of non-conforming stiffened multipatch structures. *Comput Methods Appl Mech Eng* 349:774–797. <https://doi.org/10.1016/j.cma.2019.02.042>
- Hojjat M, Stavropoulou E, Bletzinger K-U (2014) The vertex morphing method for node-based shape optimization. *Comput Methods Appl Mech Eng* 268:494–513. <https://doi.org/10.1016/j.cma.2013.10.015>
- Kloft H, Schmitz LP, Müller C, Laghi V, Babovic N, Baghdadi A (2023) Experimental application of robotic wire-and-arc additive manufacturing technique for strengthening the I-Beam profiles. *Buildings* 13(2):366. <https://doi.org/10.3390/buildings13020366>
- Misiun G, Van De Ven E, Langelaar M, Geijselaers H, Van Keulen F, Van Den Boogaard T, Ayas C (2021) Topology optimization for additive manufacturing with distortion constraints. *Comput Methods Appl Mech Eng* 386:114095. <https://doi.org/10.1016/j.cma.2021.114095>
- Mulani SB, Slemp WCH, Kapania RK (2013) EBF3PanelOpt: an optimization framework for curvilinear blade-stiffened panels. *Thin-Walled Struct* 63:13–26. <https://doi.org/10.1016/j.tws.2012.09.008>
- Sigmund O (2011) On the usefulness of non-gradient approaches in topology optimization. *Struct Multidisc Optim* 43(5):589–596. <https://doi.org/10.1007/s00158-011-0638-7>
- Sun Y, Zhou Y, Ke Z, Tian K, Wang B (2022) Stiffener layout optimization framework by isogeometric analysis-based stiffness spreading method. *Comput Methods Appl Mech Eng* 390:114348. <https://doi.org/10.1016/j.cma.2021.114348>
- Weldeyesus AG, Gondzio J, He L, Gilbert M, Shepherd P, Tyas A (2020) Truss geometry and topology optimization with global stability constraints. *Struct Multidisc Optim*. <https://doi.org/10.1007/s00158-020-02634-z>
- Zhao X, Huang H, Sun L, Chen S, Fu J (2022) Optimal design for stiffeners with a two-level approximation method involving size, shape and topology variables. *Struct Multidisc Optim* 65(9):259. <https://doi.org/10.1007/s00158-022-03371-1>
- Zhou Y, Sun Y, Xu S, Wang B (2022) Variable-height stiffener design using topology optimization with anisotropic filter-based casting constraints. *Struct Multidisc Optim* 65(11):319. <https://doi.org/10.1007/s00158-022-03428-1>
- Zienkiewicz OC, Taylor RL, Zhu JZ (2013) The finite element method: its basis and fundamentals, 7th edn. Butterworth-Heinemann, Oxford. <https://doi.org/10.1016/C2009-0-24909-9>

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