```
@inproceedings{krokowski2025penalty,
  author = {Joshua Krokowski and John T. Hwang},
  title = {Penalty-Based Coupling of Non-Matching Beam Cross-Section Meshes for
    Large-Scale Design Optimization},
  booktitle = {AIAA AVIATION FORUM AND ASCEND 2025},
  year = {2025},
  doi = {10.2514/6.2025-3473},
  url = {https://doi.org/10.2514/6.2025-3473}
}
```

Penalty-Based Coupling of Non-Matching Beam Cross-Section Meshes for Large-Scale Design Optimization

Joshua J. Krokowski*,
University of California San Diego, 9500 Gilman Dr, La Jolla, CA 92093

John T. Hwang†
University of California San Diego, 9500 Gilman Dr, La Jolla, CA 92093

This paper presents a new method for computing the cross-sectional stiffness matrix of prismatic beam structures from non-matching, overlapping finite element meshes. The approach introduces a variational coupling formulation that enforces interfacial traction compatibility and rigid body motion constraints between independently meshed subdomains using a mortar-style background mesh. A Nitsche-like term is employed to weakly enforce stress equilibrium across interfaces. This formulation enables robust handling of overlapping or non-conforming meshes without the need for remeshing or geometry pre-processing, a current challenge in multi-disciplinary design analysis and optimization. The method is compatible with natural, component-level parameterizations of structural features, such as spar cap thickness, and is designed to integrate directly into state of the art gradient-based optimization workflows critical for tasks like rotor blade shape optimization. Validation against a conformal reference solution demonstrates that the penalty-coupled approach achieves a relative error of less than 1% in the stiffness matrix while producing smooth, physically consistent warping functions for all six fundamental loading modes. The proposed method offers a flexible and optimization-ready framework for high-fidelity beam modeling in the design of advanced aerospace structures.

I. Introduction

Compared to other structural models like shells or solid elements, beam models significantly reduce the computational cost of predicting deflection, strain and stress in a wide range of aerospace structures such as high-aspect-ratio wings, rotor blades, and bracing struts, especially in dynamic simulations. In this modeling paradigm, a complex 3D structure can be represented as a topologically 1D, geometrically 3D object, without making a significant sacrifice in modeling fidelity. A beam modeling approach offloads the effort of translation between design geometry and analysis geometry to a cross-sectional analysis method, which can be formulated as a Partial Differential Equation (PDE) defined over the plane orthogonal to the 1D beam axis. This cross-sectional analysis step determines how the 3D geometry displacements, strains, and stresses relate to the 1D geometry and vice versa.

In aerospace, the design is often represented by a detailed CAD model using a boundary representation, as in the geometry handling in the aircraft design tool CADDEE[1]. Pre-processing steps such as offsetting outer surfaces or representing joined interface are a barrier to usage of high-fidelity beam models in MDO as they generally require non-differentiable operations to heal geometries prior to analysis and optimization. For rotor design, bridging the gap between these high-fidelity geometries and computation-saving beam models is critical, since the computational costs are dramatically lowered while much of the accuracy (compared to a shell or solid model) is retained.

This challenge is not present to the same degree in low-fidelity beam models that use simplified cross-sectional assumptions since the geometrical pre-processing is simplified. These lower-fidelity beam models generally use analytical expressions for structural members cross-sections, for example hollow-box or tubes. Using high-fidelity cross-sectional modeling approaches with a naturally computationally efficient beam model is a desirable target for MDO problems, as the simulation time can be significantly reduced without the loss of accuracy that comes from considering a beam to be a primitive, homogeneous geometry.

In this work, we address these challenges by extending existing work on constitutive modeling for beam structures to allow for the construction of a single set of cross-sectional properties from multiple overlapping, non-matching meshes using a penalty-based method. This approach allows for simplified construction and robust modification of

*Ph.D. Student, Department of Mechanical and Aerospace Engineering, AIAA Student Member
†Assistant Professor, Department of Mechanical and Aerospace Engineering, AIAA Senior Member

cross-sectional geometries through design iterations. Utilizing this penalty based cross-sectional analysis enables the ability to perform large-scale shape optimization for aeroelastic problems relevant to real industrial designs. The novelty of this work is to introduce a penalty-based method that eliminates the need for conformal mesh generation of beam cross-sections which enables seamless integration of overlapping, non-matching geometries in beam models for MDO. Unlike traditional methods requiring conformal mesh generation, this penalty-based approach offers scalability and flexibility for large-scale MDO problems, reducing pre-processing time without significant accuracy loss.

II. Background

Beam models are commonly used to simplify the prediction of a structure's behavior under applied loads and boundary conditions, given some kinematic assumptions. A beam model is useful when a structure is long and slender (e.g. when one among the three characteristic dimensions L, b, and h is much larger than the other two: $L \gg h, b$). If this condition is met, a beam model can reasonably simplify a full three-dimensional (3-D) solid mechanics problem to a one-dimensional (1-D) problem that depends on the position along the largest characteristic length L. Displacements, strains, and stresses at any point in the structure can be computed given only the information about the 3-D displacement (and rotation) of a "beam axis" and relevant properties of the beam's cross-section. This general approach to simplifying the 3-D solution of a solid mechanics problem is portrayed in Figure 1.

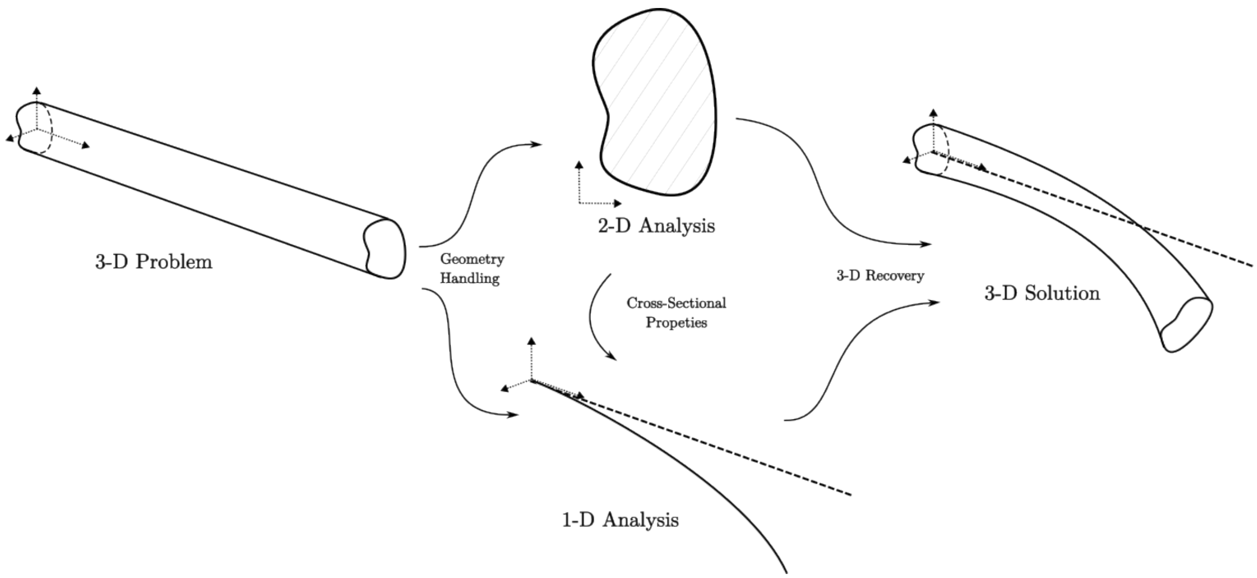

Fig. 1 Conceptual Approach for Solving 3-D Beam Problems

The primary goal of a beam model is to simplify the computation of displacements, strains, and stresses in the beam while retaining reasonable accuracy. Depending on the application, nonlinearity of kinematics and material type may need to be considered. Generally in aerospace applications, material nonlinearity is of less importance than geometric nonlinearity. For example, most common aerospace structural materials can be modeled effectively as linear elastic, such as metals or composites, with isotropic or orthotropic constitutive models respectively.

Geometric partitioning is a crucial part of any beam model. In order to carry out any beam-based analysis, a 3-D structure must be divided into 1-D and 2-D domains. This is often trivial for uniform, symmetric, isotropic, and homogeneous cross-sections, but the construction of this representation requires careful thought for beams with twists, curvatures, multiple materials, and variable cross-sections along the span. The construction of these representations is a critical step in any accurate beam analysis. First, the location (along the span) and the number of the cross-sections are selected. These cross-sections are then utilized in separate cross-sectional analysis step. Next, the spanwise direction is assigned a beam axis representation. Importantly, the beam axis construction and the geometric coordinates of the cross-section depend on each other. The beam axis properties about one point can be transformed to account for a translation of the beam axis within the cross-section. These geometries are then utilized in the analysis steps.

Using the solutions of the 1-D and 2-D analyses, the 3-D displacements, stresses, and strains are extracted. This last step is often called "3-D recovery".

A. Existing Cross-Sectional Analysis Approaches

A wide body of knowledge exists on the 2-D cross-sectional analysis step. Among the most popular are the Variational Asymptotic Beam Sectional analysis (VABS)[2] and the Beam Cross-sectional Analysis (BECAS)[3].

VABS is based on the Variational Asymptotic Method[4] and uses the strain measures to derive a set of equations that are exact up to a certain asymptotic order that can be solved in succession to compute warping functions. The warping functions can be used to describe displacements of the cross-section that occur in addition to the displacement of the cross-sections centroid. These equations are solved using the finite element method. These warping functions can then be used to compute a beam cross-sectional stiffness matrix.

BECAS uses a slightly different approach where the displacements and strain expressions are partitioned between sectional strains and axial strains. The total virtual work is derived with these partitioned strains and then a set of equilibrium equations is constructed from this total virtual work. These equilibrium equations are analyzed for 'central' and 'extremity' solutions. The 'central' solutions are those that correspond to the St.-Venant behavior of the beam section.

B. Energy-Based Polynomial Expansion

The work in this paper modifies a slightly different method outlined in a paper by Patil [5] which outlines an approach based on the total complementary potential energy. This section will review this approach and add additional detailed insights. This energy based approach is critical to the novel method implemented here and the detail in this section serves to highlight the benefits of the improvements proposed in Section III.

The method begins with a definition of the 3D displacement of a beam structure using a polynomial expression along the beam axial direction x_1 and four 3-dimensional vector displacement functions, $\overline{u}_i, \hat{u}_i, \tilde{u}_i, \breve{u}_i$:

$$u_i(x_1, x_2, x_3) = \overline{u}_i(x_2, x_3) + \hat{u}_i(x_2, x_3)x_1 + \tilde{u}_i(x_2, x_3)x_1^2 + \breve{u}_i(x_2, x_3)x_1^3\,. \tag{1}$$

The strains through the cross-section can be expressed in terms of the symmetric derivative of the displacement:

$$\epsilon_{ij} = \frac{1}{2}\left(\frac{\partial u_i}{\partial x_j} + \frac{\partial u_j}{\partial x_i}\right). \tag{2}$$

Additionally, these strains can be used to compute stresses via the material constitutive tensor, provided the material constitutive tensor is only a function of the spatial coordinate:

$$\sigma_{ij} = C_{ijkl}\epsilon_{kl}. \tag{3}$$

The static linear elastic equations with no body force applied are:

$$\sigma_{ij,j} = \left(C_{ijkl}\epsilon_{kl}\right)_{,j} = 0. \tag{4}$$

Under the assumption that the material constitutive tensor for each material used is not a function of any of the spatial coordinates, we can substitute this displacement expression into Eq. 4. Following the derivation in [5], these equations can be expressed in weak variational form posed over the beam cross-section A:

$$\begin{aligned} a^{(,0)}(\mathbf{w},\mathbf{v}) = C_{i1k1}\iint_A \tilde{u}_k\psi dA + C_{i1k\beta}\iint_A \hat{u}_{k,\beta}\psi dA - C_{i\alpha k1}\iint_A \hat{u}_k\psi_{,\alpha} dA \\ -C_{i\alpha k\beta}\iint_A \overline{u}_{k,\beta}\psi_{,\alpha} dA + \int_0^L \overline{T}_i\psi dl = 0 \end{aligned} \tag{5a}$$

$$\begin{aligned} a^{(,1)}(\mathbf{w},\mathbf{v}) = 6C_{i1k1}\iint_A \breve{u}_k\psi dA + 2C_{i1k\beta}\iint_A \tilde{u}_{k,\beta}\psi dA - 2C_{i\alpha k1}\iint_A \tilde{u}_k\psi_{,\alpha} dA - C_{i\alpha k,\beta}\iint_A \hat{u}_{k\beta}\psi_{,\alpha} dA \\ + \int_0^L \hat{T}_i\psi dl = 0 \end{aligned} \tag{5b}$$

$$a^{(,2)}(\mathbf{w},\mathbf{v}) = 3C_{i1k\beta}\iint_A \breve{u}_{k,\beta}\psi dA - 3C_{i\alpha k1}\iint_A \breve{u}_k\psi_{,\alpha} dA - C_{i\alpha k\beta}\iint_A \tilde{u}_{k,\beta}\psi_{,\alpha} dA + \int_0^L \tilde{T}_i\psi dl = 0 \tag{5c}$$

$$a^{(,3)}(\mathbf{w},\mathbf{v}) = -C_{i\alpha k\beta}\iint_A \breve{u}_{k,\beta}\psi_{,\alpha} dA + \int_0^L \breve{T}_i\psi dl = 0. \tag{5d}$$

Warping functions and their associated test functions utilize the notation,

$$\mathbf{w} = \begin{bmatrix} \overline{u}_i \\ \hat{u}_i \\ \tilde{u}_i \\ \breve{u}_i \end{bmatrix} \quad \mathbf{v} = \begin{bmatrix} \overline{\psi}_i \\ \hat{\psi}_i \\ \tilde{\psi}_i \\ \breve{\psi}_i \end{bmatrix}, \tag{6}$$

with traction-free boundary condition terms defined as

$$\begin{aligned} \overline{T}_i &= C_{i\alpha k1}\hat{u}_k n_\alpha + C_{i\alpha k\beta}\overline{u}_{k,\beta} n_\alpha \\ \hat{T}_i &= 2C_{i\alpha k1}\tilde{u}_k n_\alpha + C_{i\alpha k\beta}\hat{u}_{k,\beta} n_\alpha \\ \tilde{T}_i &= 3C_{i\alpha k1}\breve{u}_k n_\alpha + C_{i\alpha k\beta}\tilde{u}_{k,\beta} n_\alpha \\ \breve{T}_i &= C_{i\alpha k\beta}\breve{u}_{k,\beta} n_\alpha, \end{aligned} \tag{7}$$

where i, j, k, l are indices that take values $1, 2, 3$ and α, β are indices that take values $2, 3$. The notation $^{(,q)}$ signifies the associated polynomial order. A key insight is that this polynomial expansion expression from Eq. 1 must be satisfied at all x_1 and thus we get 4 different 3-Dimensional vector PDEs (one for each order of the polynomial).

Equation II.B allows us to write the weak variational form with multiple bilinear terms:

Find $\mathbf{w} = \begin{bmatrix} \overline{u}_i & \hat{u}_i & \tilde{u}_i & \breve{u}_i \end{bmatrix}^T$, $\mathbf{v} = \begin{bmatrix} \overline{\psi_i} & \hat{\psi}_i & \tilde{\psi}_i & \breve{\psi}_i \end{bmatrix}^T \in H^1(\Omega)$ with:

$$a^{(,0)}(\mathbf{w}, \mathbf{v}) + a^{(,1)}(\mathbf{w}, \mathbf{v}) + a^{(,2)}(\mathbf{w}, \mathbf{v}) + a^{(,3)}(\mathbf{w}, \mathbf{v}) = 0 \tag{8}$$

This set of equations is discretized using the finite element method and a stiffness matrix is assembled from the equations defined in Eq. :

$$\mathbf{K_{PDE}}\mathbf{w} = \mathbf{0}. \tag{9}$$

This assembled system is under-constrained, but $\mathbf{K_{PDE}}$ has a nullspace, $\mathbf{Z}$, of rank twelve. In general, finding the sparsest basis for a nullspace of a finite element problem is NP-hard, but this physical system knowledge allows us to search for only 12 singular vectors . However, there are fundamentally only twelve possible modes in the nullspace of the stiffness matrix $[K_{PDE}]$. Six modes are associated with rigid body translation and rotation. Six other modes are associated with elastic modes representing axial extension, shearing in two principal directions, torsion, and bending in two principal directions. The theory behind these null vectors is covered in the literature[6]. These 12 modes are the columns z_i of the nullspace $\mathbf{Z}$ that relate a set of solution coefficients $\mathbf{c_z}$ to the warping functions $\mathbf{w}$:

$$\mathbf{w} = \mathbf{Z}^T\mathbf{c_z} = \begin{bmatrix} | & | & & | \\ z_1 & z_2 & \dots & z_{12} \\ | & | & & | \end{bmatrix} \begin{bmatrix} c_1 \\ c_2 \\ \vdots \\ c_{12} \end{bmatrix}. \tag{10}$$

In other words, the nullspace $\mathbf{Z} \in 12 \times N_d$ forms an orthogonal basis in $\mathbb{R}^{12}$ that spans all possible displacements of the finite element mesh nodes. However, this nullspace computation is not guaranteed to find "physcially meaningful" singular vectors. Thus it is necessary to compute a transformation from the standard basis E to a reduced basis F that spans only the stress producing deformations. This computation is an orthogonalization process, so while a Gram-Schmidt style process is likely the most standard procedure, a basis transformation to a reduced nullspace $\mathbf{Z}_{red}$ can be written using a basis transformation $\mathbf{T}$:

$$\mathbf{Z}_{red} = \mathbf{TZ}, \tag{11}$$

where the nullspace $\mathbf{Z} \in \mathbb{R}^{12\times N_d}$, the reduced nullspace $\mathbf{Z}_{red} \in \mathbb{R}^{6\times N_d}$, and the transformation matrix $[T] \in \mathbb{R}^{6\times 12}$.

The original basis of the nullspace $[Z]$ is the standard basis in $\mathbb{R}^{12}$:

$$E = \{e_1, e_2, \ldots, e_{12}\} = \begin{bmatrix} 1 & 0 & & 0 \\ 0 & 1 & & 0 \\ \vdots & \vdots & \ldots & \vdots \\ 0 & 0 & & 1 \end{bmatrix} = I_{12}. \tag{12}$$

The reduced nullspace basis F is described by the functions determining the load along the span:

$$\mathcal{F} = \{f_1, f_2, \ldots, f_6\} \in \mathbb{R}^{12\times 6} = \left\{ P_1(x_1) \quad V_2(x_1) \quad V_3(x_1) \quad T_1(x_1) \quad M_2(x_1) \quad M_3(x_1) \right\}. \tag{13}$$

Any vector in the standard basis E can be represented in the new reduced basis as a linear combination of vectors in the new reduced basis F:

$$f_i = Te_i = \sum_i^{12} e_i. \tag{14}$$

The expression for the basis transformation matrix is

$$\mathbf{T} = \mathbf{F}^\dagger \mathbf{E} = \mathbf{F}^\dagger, \tag{15}$$

where $F^\dagger$ is the Moore-Penrose pseudoinverse.

The nullspace can be viewed as a collection of the warping functions expressed in the standard basis:

$$\mathbf{Z} = \begin{bmatrix} \mathbf{w}^{(1),E} \\ \mathbf{w}^{(2),E} \\ \vdots \\ \mathbf{w}^{(12),E} \end{bmatrix}. \tag{16}$$

The reduced nullspace can be seen as a collection of the warping functions expressed in the reduced basis:

$$\mathbf{Z}_{red} = \begin{bmatrix} \mathbf{w}^{(1),F} \\ \mathbf{w}^{(2),F} \\ \vdots \\ \mathbf{w}^{(6),F} \end{bmatrix}. \tag{17}$$

This gives a set of orthogonal warping functions that represent each fundamental mode. The precise displacement of the cross-section can be determined from the elastic solution coefficients via the inner product:

$$\mathbf{w} = \mathbf{Z}_{red}^T \mathbf{c}. \tag{18}$$

The basis matrix can be constructed using the functions for the loads $\mathbf{p}$ at any point along the beam cross-section, which are defined for each polynomial order $q = 0, \ldots, 3$ as:

$$\mathbf{p}^{(,q)} = \begin{bmatrix} P_1^{(,q)} \\ V_2^{(,q)} \\ V_3^{(,q)} \\ T_1^{(,q)} \\ M_2^{(,q)} \\ M_3^{(,q)} \end{bmatrix} \tag{19}$$

and the individual loads are defined as:

$$
\begin{aligned}
P_1^{(,q)}(x_2,x_3) &= \iint_\Omega \sigma_{11}^{(,q)}(x_2,x_3)\,dA \\
V_2^{(,q)}(x_2,x_3) &= \iint_\Omega \sigma_{12}^{(,q)}(x_2,x_3)\,dA \\
V_3^{(,q)}(x_2,x_3) &= \iint_\Omega \sigma_{13}^{(,q)}(x_2,x_3)\,dA \\
T_1^{(,q)}(x_2,x_3) &= \iint_\Omega \left[x_2\sigma_{13}^{(,q)}(x_2,x_3) - x_3\sigma_{12}^{(,q)}(x_1,x_2,x_3)\right]dA \\
M_2^{(,q)}(x_2,x_3) &= \iint_\Omega x_3\sigma_{11}^{(,q)}(x_2,x_3)\,dA \\
M_3^{(,q)}(x_2,x_3) &= -\iint_\Omega x_2\sigma_{11}^{(,q)}(x_2,x_3)\,dA,
\end{aligned}
\tag{20}
$$

where p, r are indices that take values $1, \ldots, 6$. Note that the stress and strain can be computed at any polynomial order of x_1, $q = 0, ..., 3$. These warping strains can be partitioned across polynomial orders:

$$\epsilon_{ij}^{(,q)} = \frac{1}{2}\left(u_{i,j}^{(,q)} + u_{j,i}^{(,q)}\right). \tag{21}$$

For $q \leq 2$:

$$u_{i,j}^{(,q)} = \frac{\partial u_i}{\partial x_j} = \begin{bmatrix} w_1^{(,q+1)} & w_{1,2}^{(,q)} & w_{1,3}^{(,q)} \\ w_2^{(,q+1)} & w_{2,2}^{(,q)} & w_{2,3}^{(,q)} \\ w_3^{(,q+1)} & w_{3,2}^{(,q)} & w_{3,3}^{(,q)} \end{bmatrix}, \tag{22}$$

and for $q = 3$:

$$u_{i,j}^{(,3)} = \frac{\partial u_i}{\partial x_j} = \begin{bmatrix} 0 & \breve{u}_{1,2} & \breve{u}_{1,3} \\ 0 & \breve{u}_{2,2} & \breve{u}_{2,3} \\ 0 & \breve{u}_{3,2} & \breve{u}_{3,3} \end{bmatrix}. \tag{23}$$

The basis transformation produces a set of six warping functions $\mathbf{w}^{(\mathbf{p})}$ that are stress-orthogonal and are weighted w.r.t. the solution coefficients $\mathbf{c_p}$ to produce a specific warping solution $\mathbf{w}$

$$\mathbf{w} = \begin{bmatrix} \mathbf{w}^{(,0)} \\ \mathbf{w}^{(,1)} \\ \mathbf{w}^{(,2)} \\ \mathbf{w}^{(,3)} \end{bmatrix} = \begin{bmatrix} \overline{u}_i \\ \hat{u}_i \\ \tilde{u}_i \\ \breve{u}_i \end{bmatrix} = \begin{bmatrix} \overline{u}_i^{(1)} & \ldots & \overline{u}_i^{(6)} \\ \hat{u}_i^{(1)} & \ldots & \hat{u}_i^{(6)} \\ \tilde{u}_i^{(1)} & \ldots & \tilde{u}_i^{(6)} \\ \breve{u}_i^{(1)} & \ldots & \breve{u}_i^{(6)} \end{bmatrix} \begin{bmatrix} c_1 \\ \vdots \\ c_6 \end{bmatrix} \tag{24}$$

where $\mathbf{w}_{i+3q}$ and q is the index of polynomial order

The elastic modes are linearly related to the reaction loads $\mathbf{p}$ at any point along the beam axis by a matrix K_1:

$$\mathbf{p} = \begin{bmatrix} P_1 \\ V_2 \\ V_3 \\ T_1 \\ M_2 \\ M_3 \end{bmatrix} = \mathbf{K_1 c}. \tag{25}$$

The matrix $\mathbf{K_1}$ is the Jacobian of the loads w.r.t. the solution coefficients $\mathbf{c}_p$:

$$\mathbf{K_1} = \frac{\partial \mathbf{p}_p}{\partial \mathbf{c}_r}. \tag{26}$$

Similarly, the total potential complementary energy is a quadratic functions of the elastic solution coefficients:

$$U^c = \frac{1}{2}\mathbf{c}^T\mathbf{K_2}\mathbf{c}. \tag{27}$$

The matrix $\mathbf{K_2}$ is the Hessian of the potential energy w.r.t. the solution coefficients $\mathbf{c_p}$:

$$\mathbf{K_2} = \frac{\partial^2 U^c}{\partial \mathbf{c}_p \partial \mathbf{c}_r}. \tag{28}$$

Both $\mathbf{K_1}$ and $\mathbf{K_2}$ can be found by taking Gateaux derivatives in the direction of the solution coefficients. Once these two matrices are computed, we have the following relationships for the beam constitutive matrix $\mathbf{K}$and the beam compliance matrix $\mathbf{S}$ in terms of $\mathbf{K_1}$ and $\mathbf{K_2}$:

$$\begin{aligned} \mathbf{S} &= \mathbf{K_1}^{-T}\mathbf{K_2}\mathbf{K_1}^{-1} \\ \mathbf{K} = \mathbf{S}^{-1} &= \left(\mathbf{K_1}^{-T}\mathbf{K_2}\mathbf{K_1}^{-1}\right)^{-1} \end{aligned} \tag{29}$$

III. Beam Cross-Sectional Modeling with Non-Conformal, Overlapping Meshes

A. Conformal Cross-Sectional Analysis

This section covers the modifications to the energy-based expansion approach outlined in Section II.B aimed at handling multiple, non-matching, overlapping meshes in the context of gradient-based design optimization. Instead of using a nullspace computation and a basis transformation to find a set of of load-orthogonal warping functions, we apply a set of constraints to the variational form and enforce the constraints via Lagrange multipliers. Assembly of the system matrix is completed and the linear system is solved with 6 different right hand side vectors, each one serving to produce a separate warping mode.

1. Lagrange Multiplier Enforced Constraints

In this section, we define a constraint matrix $C(\mathbf{w}) = F_c \in \mathbb{R}^{30}$ via Lagrange multipliers $\lambda \in \mathbb{R}^{30}$. These constraints consist of:

- **6 rigid body constraints**, enforcing orthogonality of $\bar{u}$ to rigid motion by setting the rigid rotations and displacements to be equal to 0:

$$\begin{aligned} C_i^{\text{translation}}(\mathbf{w}) &= \int_\Omega \bar{u}_i \, d\Omega = 0 \qquad \text{for } i = 1,2,3 \quad \text{(translation)} \\ C_i^{\text{rotation}}(\mathbf{w}) &= \int_\Omega \varepsilon_{ijk} x_\alpha \bar{u}_k \, d\Omega = 0 \quad \text{for } i = 1,2,3 \quad \text{(rotation)} \end{aligned}$$

- **24 Loading Constraints** enforce that the warping functions that produce the stresses that produce a specific force over the cross-section are orthogonal to each other and of unit magnitude:

$$C_{p+6q}^{\text{load}}(\mathbf{w}) = \int_\Omega \mathbf{p}^{(p,q)} \, d\Omega = f_{p+6q} \quad \text{for } (p = 1,\ldots,6), (q = 0,\ldots,3) \quad \text{(loading)}$$

  where $f_{p+6q} = 0$ unless $r = p$. In this case, r is the index corresponding to a specific warping solution (e.g. a "load case) and p is the generalized stress/strain mode.

The assembly of the system stiffness matrix $\mathbf{K_{PDE}}$ remains the same as Eq. 9, but now we have augmented the system with the constraints:

$$\begin{aligned} &\sum_{q=0}^{3} a^{(,q)}(\mathbf{w},\mathbf{v}) + \sum_{i=1}^{24}\lambda_i C_i^{\text{load}}(\mathbf{w}) + \sum_{25}^{27}\lambda_i C_i^{\text{translation}}(\mathbf{w}) + \sum_{i=28}^{30}\lambda_i C_i^{\text{rotation}}(\mathbf{w}) \\ &\quad + \sum_{i=1}^{24}\mu_i C_i^{\text{load}}(\mathbf{v}) + \sum_{25}^{27}\mu_i C_i^{\text{translation}}(\mathbf{v}) + \sum_{i=28}^{30}\mu_i C_i^{\text{rotation}}(\mathbf{v}) \qquad = \sum_{i=1}^{30}\mu_i F_i. \end{aligned} \tag{30}$$

This weak form results in a discrete KKT system to solve :

$$\begin{bmatrix} K_{PDE} & C^\top \\ C & 0 \end{bmatrix} \begin{bmatrix} \mathbf{w} \\ \lambda \end{bmatrix} = \begin{bmatrix} F \\ F_c \end{bmatrix}. \tag{31}$$

2. Warping Functions from Multiple Right Hand Side Solves

The resultant system from Eq. 31 must be solved six times with different right hand side (RHS) vectors, but the same left hand side (LHS) operator can be reused. The first 6 entries of F_c are the only values to change throughout these 6 solves. This approach saves both computational effort and memory overhead, compared to the nullspace based approach or other eigenvalue based approaches from literature.

3. Beam Cross-Sectional Stiffness Matrix Computation

The procedure for computing the beam cross-sectional stiffness matrix remains identical to that presented in Section II.B. The implementation of the constrained approach is to simplify the computation of the warping functions, but the procedure for assembling the beam cross-sectional stiffness matrix remains unchanged.

B. Penalty-Based Coupling of Overlapping, Nonconforming Cross-Sectional Meshes

1. Overlapping Domains

In the most general nonconformal setting, the cross-section of a beam may be discretized into multiple independently generated subdomains that geometrically overlap without perfectly matching boundaries. As a result, the union of these subdomain meshes does not exactly represent the physical overlapping region. To resolve this, we adopt a penalty-based mortar method in which a *background mesh* is constructed to cover the overlapping region. While mortar methods are commonly used along with lagrange multipliers[7], there are examples of implementations in mechanics of penalty-based mortar methods such as the isogeometric shell coupling code PenGoLINS[8]. A background mesh serves as a consistent integration domain and facilitates variational coupling across sub-domain interfaces. The mortar method is popular in solid mechanics, especially so in contact mechanics, where these types of saddle point problems routine arise[9]

Although the background mesh is necessarily non-conformal to either sub-mesh, interpolation or projection can be used to transfer functions on one of the sub-domains to the mortar mesh. The mortar mesh enables accurate numerical integration over the true geometric overlap. The resulting geometric approximation error is typically small and can be reduced systematically through refinement. This error is generally dominated by the discretization error within each subdomain.

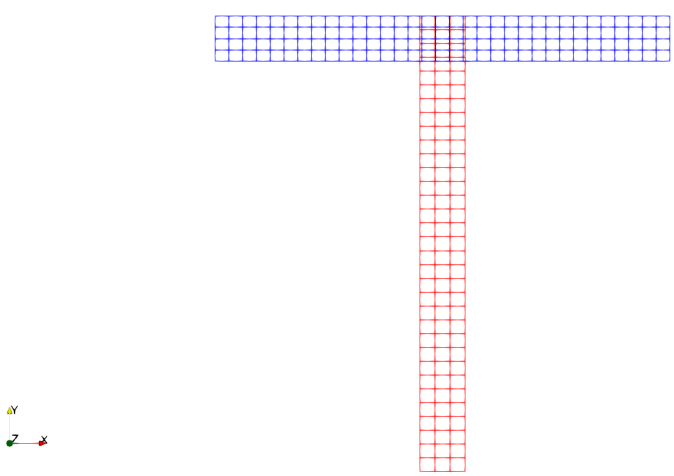

Fig. 2 Non-matching, Overlapping 2D Meshes

2. Coupling of Warping Functions

The computation of the beam's cross-sectional stiffness matrix in the non-conformal case is primarily about coupling the warping functions defined over each subdomain. The core idea is to extend the conformal formulation by adding variational terms that enforce displacement and traction compatibility across subdomains.

Specifically, we augment the weak form with two penalty-based coupling terms:

- A **displacement penalty** term, which penalizes discontinuities in the warping displacement field across the interface,
- A **traction penalty** term, which penalizes mismatched tractions on the immersed boundaries of overlapping subdomains.

Both terms are integrated over the background mesh. To couple fields defined on separate subdomains, projection operators are constructed to map degrees of freedom from each subdomain onto the background mesh. These projection operators are built using only the *colliding cells*—those elements of each mesh that intersect the background domain—thus greatly reducing computational cost.

To describe this approach, consider the coupling of two non-matching finite element discretizations, defined on overlapping beam cross-sectional subdomains. Let Ω_A and Ω_B denote two subdomains of the cross-section, with an overlapping region $\Omega^C = \Omega_A \cap \Omega_B$. Each subdomain is meshed independently, giving rise to finite element spaces V_A and V_B, with associated warping functions $\mathbf{w}_A \in V_A$, $\mathbf{w}_B \in V_B$.

The unconstrained PDE from Equation II.B is posed separately on each subdomain: Find $\mathbf{w} = \begin{bmatrix} \overline{u}_i & \hat{u}_i & \tilde{u}_i & \breve{u}_i \end{bmatrix}^T$, $\mathbf{v} = \begin{bmatrix} \overline{\psi}_i & \hat{\psi}_i & \tilde{\psi}_i & \breve{\psi}_i \end{bmatrix}^T \in H^1(\Omega)$ with:

$$\sum_{q=0}^{3} a^{(,q)}(\mathbf{w}_A, \mathbf{v}_A) + \sum_{q=0}^{3} a^{(,q)}(\mathbf{w}_B, \mathbf{v}_B) = 0 \tag{32}$$

To enforce compatibility between $\mathbf{w}_A$ and $\mathbf{w}_A$ within the overlapping region, we define a third mesh $\mathcal{T}^C$ over Ω^C, typically constructed from the union of overlapping cells in $\mathcal{T}_A$ and $\mathcal{T}_B$, with boundary intersection refinement as needed. This mesh must be constructed as the union of the two overlapping meshes, see Figure 3 for an example of this process.

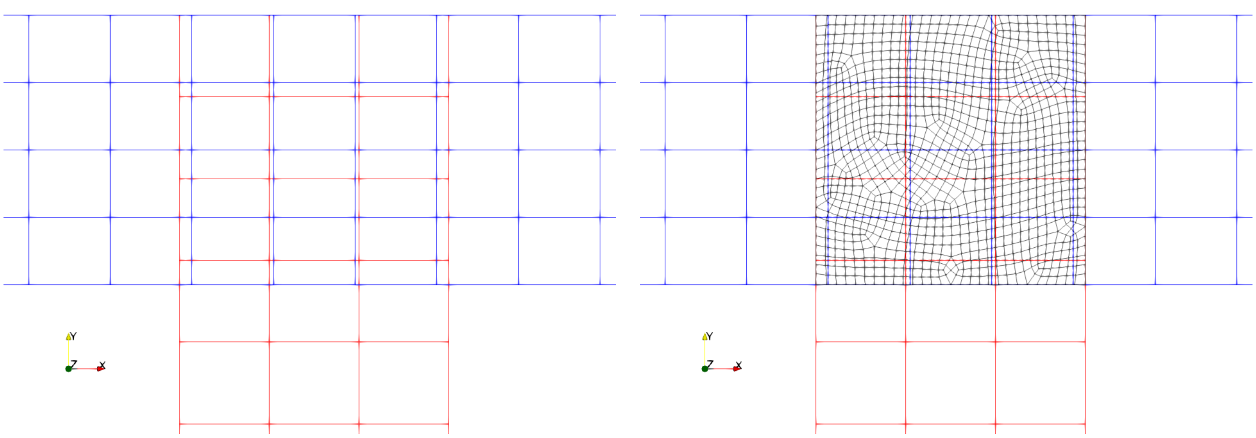

(a) Non-conformal meshes with overlap **(b) Automatically Generated Background Mesh**

Fig. 3 Mortar Mesh Construction

We introduce projection (interpolation) operators from each original space to the overlapping background space:

$$\mathbf{P}_A : V_A \to V_C, \qquad \mathbf{P}_B : V_B \to V_C, \tag{33}$$

which map functions defined on $\mathcal{T}_A$ and $\mathcal{T}_B$ into the common function space V_C over $\mathcal{T}^C$.

A volumetric penalty term is defined to weakly enforce continuity of the warping displacement field:

$$a_u(\mathbf{w}_A, \mathbf{w}_B; \mathbf{v}_A, \mathbf{v}_B) = \int_{\Omega_C} \frac{\eta_u}{h} \left(\mathbf{P}_A \mathbf{w}_A - \mathbf{P}_B \mathbf{w}_B\right) \cdot \left(\mathbf{P}_A \mathbf{v}_A - \mathbf{P}_B \mathbf{v}_B\right) \, d\Omega_C, \tag{34}$$

where $\eta_u > 0$ is a user-defined penalty constant, and h is a representative mesh size associated with $\mathcal{T}^C$.

This formulation ensures that coupling is enforced only in the physically overlapping region Ω^C and introduces no constraints or artificial stiffness outside the overlap. The use of projection operators provides smooth adjoint and shape sensitivity propagation, and the entire method fits naturally into variational solvers using FEniCS or similar finite element frameworks.

The variational problem then becomes: Find $(\mathbf{w}_A, \mathbf{w}_B) \in V_A \times V_B$ such that:

$$a_A(\mathbf{w}_A, \mathbf{v}_A) + a_B(\mathbf{w}_B, \mathbf{v}_B) + a_u(\mathbf{w}_A, \mathbf{w}_B; \mathbf{v}_A, \mathbf{v}_B) = 0 \quad \forall (\mathbf{v}_A, \mathbf{v}_B) \in V_A \times V_B. \tag{35}$$

To enforce strain compatibility across the overlap Ω^c for the constant polynomial terms of order $q = 0$, we use the previously defined projection operators $\mathbf{P}_A$ and $\mathbf{P}_B$

The traction penalty term is defined as

$$a_t(\mathbf{w}_A, \mathbf{w}_B; \mathbf{v}_A, \mathbf{v}_B) = \int_{\Gamma_C} \frac{\eta_t}{h} \left[\left(\sigma_{ij}^{(,0)} \left(\mathbf{P}_A \mathbf{w}_A\right) n_j - \sigma_{ij}^{(,0)} \left(\mathbf{P}_B \mathbf{w}_B\right) n_j \right) \cdot \left(\sigma_{ij}^{(,0)} \left(\mathbf{P}_A \mathbf{v}_A\right) n_j - \sigma_{ij}^{(,0)} \left(\mathbf{P}_B \mathbf{v}_B\right) n_j \right) \right] d\Gamma_C. \tag{36}$$

These penalty terms can be constructed by simply integrating these forms over Ω_C using the respective trial and test functions. This allows us to avoid incorporating any new degrees of freedom into the variational problem.

The penalty-weighted mass matrix on the background mesh is computed and stored:

$$M_C = \int_{\Omega} \frac{\eta_u}{h} \mathbf{w}_C \cdot \mathbf{v}_C d\Omega, \tag{37}$$

and the traction-penalty matrix is computed and stored as well:

$$S_C = \int_{\Gamma_C} \frac{\eta_u}{h} \left(\sigma_{ij}^{(,0)} (\mathbf{w}_C) n_j \right) \cdot \left(\sigma_{ij}^{(,0)} (\mathbf{v}_C) n_j \right) \, d\Gamma. \tag{38}$$

It is worth noting that, although the background mesh is utilized to assemble these two matrices, there is never a linear system solved that includes the degrees of freedom of $\mathbf{w_c}$. The background mesh construction is only there to provide an accurate representation of the overlapping area.

3. Constraint Preservation in the Coupled System

A critical property of the penalty-coupled approach is that the number of global constraints—enforced via Lagrange multipliers—remains unchanged. This includes both the six rigid-body motion constraints and the six generalized stress/moment constraints used to isolate modal warping fields. Despite the decomposition of the domain into overlapping subdomains, the physical behavior of the beam cross-section is preserved, and no additional multipliers are introduced.

The overall KKT system for two coupled domains can be constructed as:

$$\begin{bmatrix} K_A + \mathbf{P}_A^\top (M_C + S_C) \mathbf{P}_A & -\mathbf{P}_A^\top (M_C + S_C) \mathbf{P}_B & C_A^\top \\ -\mathbf{P}_B^\top (M_C + S_C) \mathbf{P}_A & K_B + \mathbf{P}_B^\top (M_C + S_C) \mathbf{P}_B & C_B^\top \\ C_A & C_B & 0 \end{bmatrix} \begin{bmatrix} \mathbf{w}_A \\ \mathbf{w}_B \\ \lambda \end{bmatrix} = \begin{bmatrix} F_A \\ F_B \\ F_c \end{bmatrix}, \tag{39}$$

Here:

- K_A, K_B are the stiffness matrices assembled on $\mathcal{T}_A$ and $\mathcal{T}_B$,
- M_C is the mass matrix of Ω_C ,
- S_C is an assembled stiffness matrix of Ω_C ,
- C_A, C_B are the constraint matrices from V_A and V_B,
- λ are the Lagrange multipliers enforcing the global constraints.
- F_A, F_B are the body force vectors for mesh A and B
- F_c is the vectors resultant values of lagrange multiplier equations.

This formulation enforces displacement compatibility in the overlapping region while respecting global physical and rigid-body constraints, and is suitable for adjoint-based sensitivity analysis and optimization over non-matching subdomains.

Each global Lagrange multiplier is mapped back to the relevant degrees of freedom on the corresponding subdomain(s), ensuring that the constraints are applied consistently and without over-constraining the system.

4. Constitutive Matrix Assembly

Following the conformal formulation, the beam constitutive matrix $\mathbf{K}$ is computed from intermediate matrices $\mathbf{K}_1$ and $\mathbf{K}_2$, which represent virtual work and constraint enforcement, respectively. In the non-conformal case, these matrices are assembled as additive contributions from each subdomain:

$$\mathbf{K}_1 = \sum_m \mathbf{K}_1^{(m)}, \qquad \mathbf{K}_2 = \sum_m \mathbf{K}_2^{(m)}, \tag{40}$$

where the superscript (m) denotes the contribution from subdomain m.

$$\mathbf{K} = \mathbf{K}_1^\top \mathbf{K}_2^{-1} \mathbf{K}_1. \tag{41}$$

5. Solution Recovery and Post-Processing

Because each unit load case is solved independently, solution recovery and post-processing can be performed at the subdomain level. The specific warping solution $\mathbf{w}$ and the stored subdomain-level contributions to $\mathbf{K}_1$ contain all the necessary information to reconstruct the stress, strain, and displacement fields within each subdomain.

This modular structure also supports local post-processing for stress recovery, shape sensitivity analysis, and evaluation of design quantities tied to specific geometric features—particularly valuable when subdomains are tied to distinct structural components such as spar caps or skin panels.

IV. Software Implementation

This section outlines the key components of the software implementation used to support the penalty-based coupling of nonconforming, overlapping cross-sectional meshes. The implementation is modular, parallelizable, and designed to accommodate both conforming and nonconforming mesh configurations, including arbitrary number of meshes.

A. Overview

The implementation in ALBATROSS[10] consists of a set of routines for detecting mesh collisions, constructing background meshes, assembling coupling operators, and computing the augmented variational forms. The core algorithm follows a three-stage process: (1) identifying overlapping regions, (2) constructing the background mesh, and (3) assembling penalty terms into the global weak form. The software is written using FEniCSx[11] and contained within a *CoupledXSProblem()* python class.

B. Collision Detection

To enable coupling between independently meshed subdomains, we first identify where those domains geometrically intersect. This is achieved using a built in FEniCSx tools for collision detection based on boundary box trees. This preprocessing steps helps future work and restrict the construction of the interpolation operators.

1. Collisions

Each sub-mesh is spatially indexed using a bounding box trees. We loop over each mesh and detect and label the cells that collide with another mesh. We also construct an adjacency matrix that stores the presence of intersections between meshes.

2. Separations

If no collision is detected between two meshes, a *separation* is recorded between those too meshes. This simply means that the penalty terms are not attempted to be constructed between these pairs of meshes and placeholder "off-diagonal blocks" are stored for the assembly of the global nested system. A separation between two regions means these two regions are treated as uncoupled and their warping functions have no direct dependence on each other.

3. Regions

A region is a portion of the global mesh (e.g. one of the sub-meshes). Each region's cells are tagged as one of three types:

- **Fully overlapping** mesh regions are cells that have all dofs fully contained within the domain of the overlapping mesh,
- **Partially overlapping** mesh regions are cells that have all dofs fully contained within the domain of the overlapping mesh,
- **Non-Overlapping** mesh regions are cells that have no dofs contained within the domain of the overlapping mesh.

These classifications are used to determine where penalty terms should be applied and which interpolation matrices need to be constructed.

C. Penalty Parameter Selection

Penalty parameters for displacement and traction compatibility terms are selected based on the characteristic stiffness and element size of the local mesh. The default penalty parameter values are:

$$\frac{\eta_u}{h} = \frac{1e7}{h}$$
$$\frac{\eta_u}{h} = 1$$

D. Accommodating Overlap

Geometric overlap between sub-meshes is accommodated through the use of a *background mortar mesh*. This mesh is constructed over the union of all detected colliding regions and is used to define integration domains for the penalty terms. Fields from each sub-mesh are interpolated onto the background mesh to enable weak enforcement of interface compatibility. The background mesh is built adaptively based on the resolution of the colliding cells, and is typically much smaller than either subdomain mesh.

Additionally, in to the overlapping section the local stiffness is "smeared" between the two overlapping meshes. For isotropic materials, this is simply adjusting the total stiffness by $1/\sqrt{n}$ meshes intersecting.

E. Interpolation Matrices

For each pair of colliding subdomain cells, an interpolation matrix is constructed that maps basis functions from the local sub-mesh onto the quadrature points of the background mesh. These matrices are sparse and localized, constructed only for intersecting regions. The mapping is used both to interpolate displacements and compute projected tractions. The interpolation operator for mesh A is denoted $\mathbf{P}_A$, and the projection between meshes is defined as:

$$\mathbf{P}_A^\top \mathbf{M} \mathbf{P}_A$$

where $\mathbf{M}$ is a mass or mass-like matrix defined over the background mesh.

F. Penalty Term Construction

The final system includes two interface penalty terms:

- **Displacement penalty term**: weakly enforces displacement continuity across subdomain interfaces,
- **Traction penalty term**: weakly enforces balance of normal tractions across subdomain interfaces.

The assembled system is symmetric, sparse, and compatible with automatic differentiation and adjoint-based sensitivity analysis frameworks.

V. Verification and Validation

A. Isotropic T-Section Validation

Fig. 4 shows the comparison of a T-section mesh used to find the cross-sectional properties of a beam. Tables 1-4 show the resulting stiffnesses of this section, given isotropic properties of E=100 and ν=0.2.

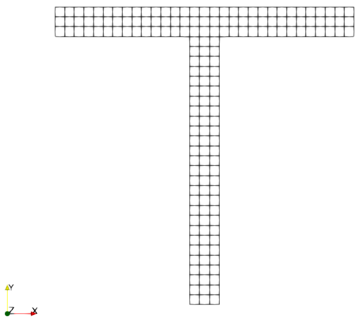

(a) Conformal Mesh

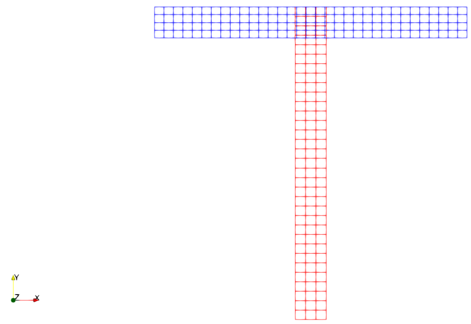

(b) Non-matching, Overlapping 2D Meshes

Fig. 4 Two Representations of a T-Section

	Tension	Shear 1	Shear 2	Torsion	Bending 1	Bending 2
Tension	1.900e+01	0	0	0	4.050e+00	2.506e-14
Shear 1	0	3.657e+00	5.779e-13	-1.615e+00	0	0
Shear 2	0	5.779e-13	3.287e+00	2.289e-13	0	0
Torsion	0	-1.615e+00	2.289e-13	7.394e-01	0	0
Bending 1	4.050e+00	0	0	0	2.663e+00	-1.920e-14
Bending 2	2.506e-14	0	0	0	-1.920e-14	8.409e-01

Table 1 Conformal Beam Constitutive Matrix

	Tension	Shear 1	Shear 2	Torsion	Bending 1	Bending 2
Tension	1.894e+01	0	0	0	4.029e+00	-2.801e-09
Shear 1	0	3.717e+00	6.193e-09	-1.632e+00	0	0
Shear 2	0	6.193e-09	3.487e+00	1.340e-08	0	0
Torsion	0	-1.632e+00	1.340e-08	7.442e-01	0	0
Bending 1	4.029e+00	0	0	0	2.656e+00	-1.170e-09
Bending 2	-2.801e-09	0	0	0	-1.170e-09	8.409e-01

Table 2 Non-conformal Beam Constitutive Matrix

	Tension	Shear 1	Shear 2	Torsion	Bending 1	Bending 2
Tension	-5.584e-02	0	0	00	-2.116e-02	-2.801e-09
Shear 1	0	6.031e-02	6.193e-09	-1.714e-02	0	0
Shear 2	0	6.193e-09	1.788e-01	1.340e-08	0	0
Torsion	0	-1.714e-02	1.340e-08	4.847e-03	0	0
Bending 1	-2.116e-02	0	0	0	-7.892e-03	-1.170e-09
Bending 2	-2.801e-09	0	0	0	-1.170e-09	-5.233e-05

Table 3 Absolute Difference, Non-matching vs. Conformal Beam Constitutive Matrix

	Tension	Shear 1	Shear 2	Torsion	Bending 1	Bending 2
Tension	2.939e-03	0	0	0	5.224e-03	1.118e+05
Shear 1	0	1.649e-02	1.072e+04	1.061e-02	0	0
Shear 2	0	1.072e+04	5.439e-02	5.855e+04	0	0
Torsion	0	1.061e-02	5.855e+04	6.555e-03	0	0
Bending 1	5.224e-03	0	0	0	2.963e-03	6.095e+04
Bending 2	1.118e+05	0	0	0	6.095e+04	6.223e-05

Table 4 Relative Difference, Non-matching vs. Conformal Beam Constitutive Matrix

The visual inspection of the warping functions shows nearly identical agreement between the penalty-coupled approach and the non-conformal approach.

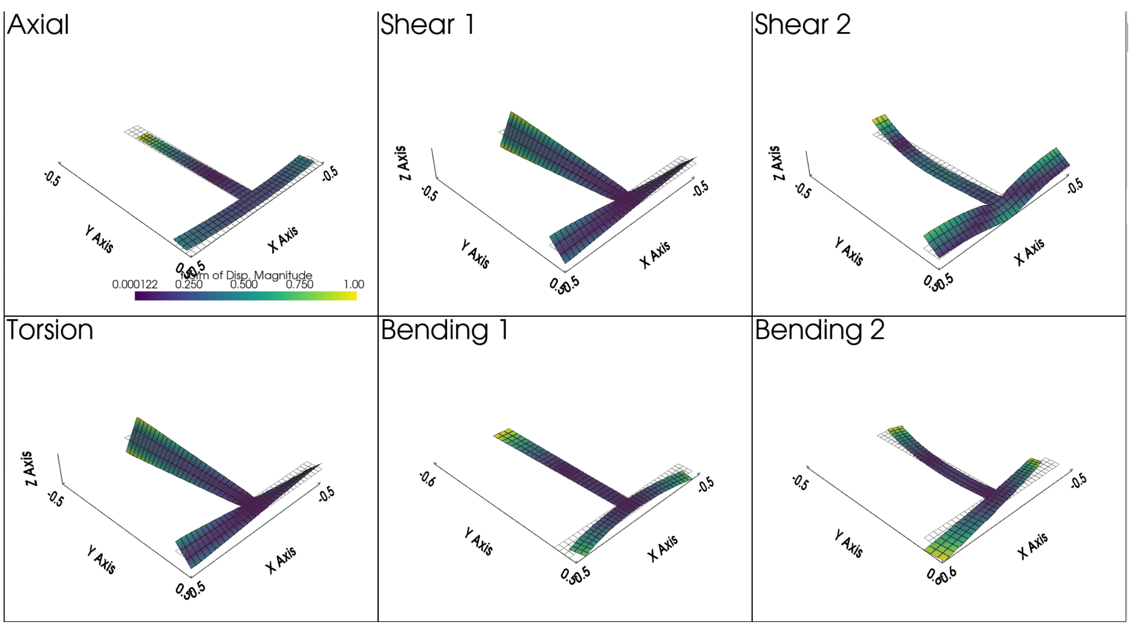

Fig. 5 Conformal Mesh Warping Functions for T-Section

The model shows high accuracy across nearly 7 orders of magnitude for the displacement based penalty parameter

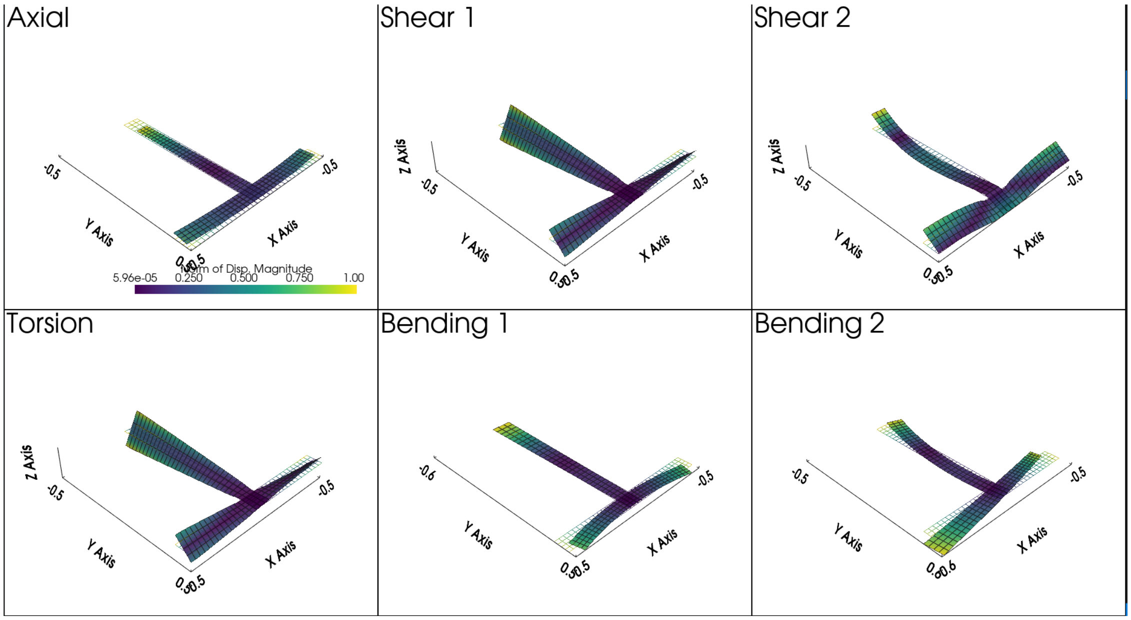

Fig. 6 Non-matching, Overlapping Meshes Warping Functions for T-Section:

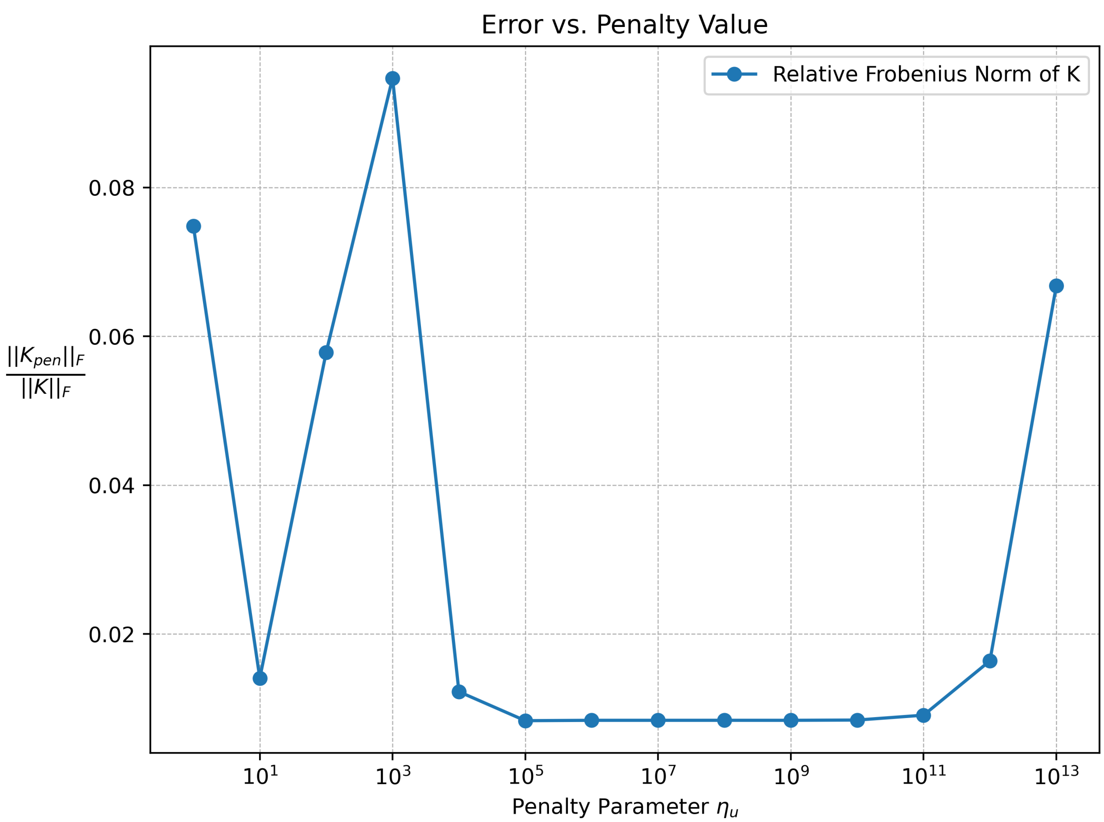

Fig. 7 Penalty Parameter Values vs Error (N=13 elements through the flange and web thickness)

The model shows healthy convergence trend, with overall error of less than 1% fairly limited number of elements. This is encouraging and suggests this method has minimal downside for a highly flexible approach.

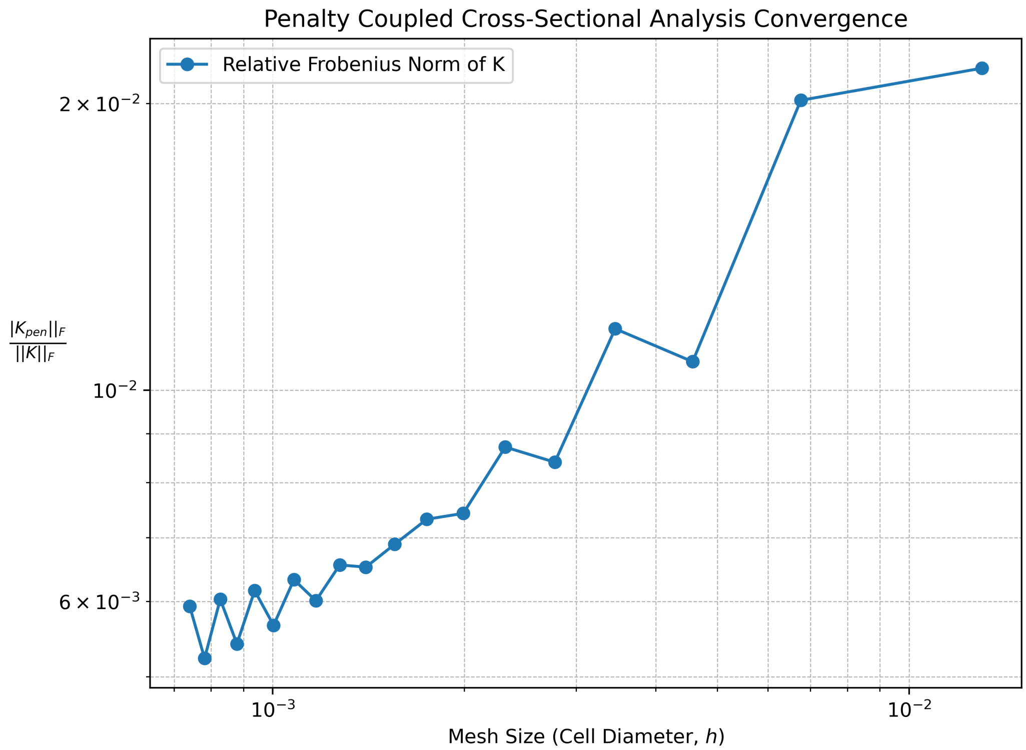

Fig. 8 Mesh Convergence Towards Conformal Solution

VI. Conclusion

This work presents a penalty-based coupling method for computing beam cross-sectional stiffness matrices using non-matching, overlapping sub-meshes. By introducing a mortar-style background mesh and augmenting the variational formulation with both a displacement continuity penalty and a Nitsche-style traction compatibility term, the method enforces physical consistency across subdomain interfaces without requiring mesh conformity.

The proposed approach achieves a stiffness matrix within 1% error of the corresponding conformal solution, demonstrating both numerical accuracy and robustness. Importantly, this formulation enables modular, component-level parameterizations—such as independent spar cap or skin thicknesses—that are difficult to accommodate in traditional conforming mesh frameworks. These capabilities are well suited to future work in the gradient-based optimization workflows critical for large-scale multi-disciplinary design optimization. These features such as avoiding complex remeshing or artificial geometry merging are important for future beam shape optimization problems that are of interest to industry.

Overall, the method enables accurate, robust, and differentiable modeling of beam cross sections without requiring a globally conformal mesh. By leveraging variational interface coupling and automated background mesh construction, this framework supports complex geometric configurations and broadens the feasible design space for cross-sectional optimization.

Looking forward, this approach provides a foundation for efficient shape and thickness optimization of slender structures. Its compatibility with non-conformal, modular geometries makes it particularly attractive for applications such as rotor blade design, where detailed structural features vary independently. The next step is to integrate this formulation into a full beam solver and gradient-based optimization loop, enabling end-to-end optimization of high-performance aerospace structures.

Acknowledgments

The work in this paper was supported by NASA under award No. 80NSSC21M0070. The author gratefully acknowledges Professor Nicholas Boechler (UC San Diego) and Professor Darshan Sarojini for their valuable support and guidance throughout this work. The author also acknowledges the use of OpenAI's ChatGPT for assistance in refining technical writing and clarifying implementation details during the preparation of this manuscript.

References

[1] Ruh, M. L., Sarojini, D., Fletcher, A., Asher, I., and Hwang, J. T., "Large-scale multidisciplinary design optimization of the NASA lift-plus-cruise concept using a novel aircraft design framework," *arXiv preprint arXiv:2304.14889*, 2023.

[2] Yu, W., Hodges, D. H., and Ho, J. C., "Variational asymptotic beam sectional analysis–an updated version," *International Journal of Engineering Science*, Vol. 59, 2012, pp. 40–64.

[3] Blasques, J. P., "User's manual for BECAS. A cross section analysis tool for anisotropic and inhomogeneous beam sections of arbitrary geometry," Tech. rep., Technical Univ. of Denmark, 2012.

[4] Berdichevsky, V., "Variational-asymptotic method of shell theory construction," *PMM*, Vol. 43, No. 4, 1979, pp. 664–687.

[5] Patil, M., "Prof. Hodges Memorial Session: A New Approach to Beam Cross-sectional Analysis," *AIAA SCITECH*, 2023. https://doi.org/10.2514/6.2023-0766, URL https://arc.aiaa.org/doi/abs/10.2514/6.2023-0766.

[6] Morandini, M., Chierichetti, M., and Mantegazza, P., "Characteristic behavior of prismatic anisotropic beam via generalized eigenvectors," *International Journal of Solids and Structures*, Vol. 47, No. 10, 2010, pp. 1327–1337.

[7] Belgacem, F. B., "The mortar finite element method with Lagrange multipliers," *Numerische Mathematik*, Vol. 84, 1999, pp. 173–197.

[8] Zhao, H., Liu, X., Fletcher, A. H., Xiang, R., Hwang, J. T., and Kamensky, D., "An open-source framework for coupling non-matching isogeometric shells with application to aerospace structures," *Computers & Mathematics with Applications*, Vol. 111, 2022, pp. 109–123.

[9] Svenning, E., "A weak penalty formulation remedying traction oscillations in interface elements," *Computer Methods in Applied Mechanics and Engineering*, Vol. 310, 2016, pp. 460–474.

[10] Krokowski, Joshua, "ALBATROSS: Analysis Library for Beams amd Three-dimensional Representations of Slender Structures," , 2024. Accessed November 21, 2024. https://github.com/jkrokowski/ALBATROSS.

[11] "FEniCS Project," , 2023. Accessed March 24, 2023. https://fenicsproject.org/.