EXPLORING MULTI-FIDELITY AEROELASTIC TAILORING: PROSPECT AND MODEL ASSESSMENT

Hauke Maathuis[1], Saullo G.P. Castro[2] & Roeland De Breuker[3]

[1]Ph.D. Candidate, Faculty of Aerospace Engineering, Delft University of Technology, h.f.maathuis@tudelft.nl
[2]Associate Professor, Faculty of Aerospace Engineering, Delft University of Technology, s.g.p.castro@tudelft.nl
[3]Associate Professor, Faculty of Aerospace Engineering, Delft University of Technology, r.debreuker@tudelft.nl

Abstract

The design and optimisation of aircraft wings are critical tasks in aerospace engineering, requiring a balance between structural integrity, aerostructural performance, and manufacturability. This multifaceted challenge involves the interplay of various disciplines, each with distinct parameters and constraints. Traditional design approaches often fall short, necessitating advanced methodologies like Multidisciplinary Design Optimisation (MDO). MDO integrates aerodynamic, structural, and manufacturability analyses to explore a vast design space and identify optimal solutions that meet performance, safety, and cost criteria.

Advancements in manufacturing technologies and material sciences have led to the increased use of composite materials, which offer an excellent weight-to-strength ratio. Aeroelastic Tailoring, which incorporates directional stiffness into structural design, further enhances performance. This study employs lamination parameters to efficiently represent composite layups within a gradient-based optimisation process, aiming to minimise weight while ensuring feasibility across multiple constraints.

The work highlights the challenge of optimising aircraft designs using multiple models of varying fidelity. Traditional sequential optimisation approaches, which progressively integrate disciplines, may miss potential superior designs due to limited initial information. Instead, concurrent optimisation schemes are explored, utilising both low-fidelity (beam-based) and high-fidelity (shell-based) models. This approach promises structural feasibility, reduces computational costs, and incorporates high-fidelity information early in the design process.

The envisioned methodology bridges different design stages, enabling better overall aircraft performance. By aligning and comparing a beam-based and shell-based model, the study explores their use in multi-fidelity optimisation. The results demonstrate the feasibility and benefits of this approach, offering a robust framework for future aircraft design projects.

Keywords: Multi-Fidelity Optimisation, Aeroelastic Tailoring, Multidisciplinary Design Optimisation, Preliminary Aircraft Design

1. Introduction

The design and optimisation of aircraft wings stand at the forefront of aerospace engineering, demanding a meticulous balance between structural integrity, aerostructural performance, and manufacturability. The complexity of this task is amplified by the interplay of multiple disciplines, each governed by its own set of parameters and constraints. Traditional design approaches often fall short in addressing this multi-faceted challenge, necessitating the adoption of advanced methodologies that can seamlessly integrate diverse disciplinary insights.

The increase in performance, or in other words a weight reduction of the structure, leads to lighter and slender wings which pose challenges like low natural frequencies coupling with aircraft motions or structural failures like buckling, due to thin-walled design. Early consideration of these disciplines in aircraft development using Multidisciplinary Design Optimisation (MDO) ensures an optimal yet safe structure. MDO emerges as a powerful framework to tackle this complexity, offering a holistic

approach that simultaneously considers the myriad factors influencing wing design. By integrating aerodynamic, structural, and manufacturability analyses, MDO enables the exploration of a vast design space to identify solutions that meet performance, safety, and cost criteria.

Thanks to the advancements in manufacturing technologies and material sciences, composite materials are being used increasingly more due to their excellent weight-to-strength ratio. A promising technology to improve the performance of aircraft by optimising the structure is the application of Aeroelastic Tailoring. Thereby, directional stiffness is embodied into the structural design to control aeroelastic deformation to beneficially impact on the aerodynamic and structural performance of the aircraft [1]. Within this work, the concept of lamination parameters provides a robust means of representing composite layups. Lamination parameters abstract the detailed layup configuration into a set of variables that can be efficiently handled within a (gradient-based) optimisation process. By representing the stiffness and strength characteristics of composite materials in a concise form, lamination parameters enable the precise tailoring of material properties to meet specific design requirements while lowering the number of design variables to be considered by the optimiser. However, the use of lamination parameters demands a subsequent discrete optimisation to retrieve stacking sequences [2].

Thus, in the following work, the emerging mathematical problem can be summarised as finding a combination of lamination parameter and thicknesses of the panels to obtain the minimum weight while ensuring feasibility, evaluated based on G multidisciplinary constraints, formulated as

$$\min_{\mathbf{x} 2 X} f(\mathbf{x}) \text{ s.t.} 8i \; 2 \; f1; :::; Gg; c_i(\mathbf{x}) \quad 0 \tag{1}$$

where $f(\mathbf{x})$ represents the weight as the objective function, $\mathbf{x} = f\mathbf{x}_{lp}; \mathbf{x}_t g \; 2 \; X \quad \mathrm{R}^D$ is the D-dimensional vector of design variables within the feasible design space, consisting of $\mathbf{x}_{lp}$ as the design variables representing lamination parameters and $\mathbf{x}_t$ the thicknesses.

Typically in aircraft design, multiple models exist, each relying on different formulations, such as beam or shell formulations for structural analysis and panel codes or CFD analysis for aerodynamics. The process of aircraft development usually begins with a pre-conceptual design from which an initial configuration is derived. Following this, preliminary sizing is conducted using a simplified model (hereafter referred to as the low-fidelity model). Even at this early stage, it is desirable to involve as many disciplines as possible to achieve an optimal design. However, this is often computationally infeasible. Instead, current design processes employ a sequential approach where disciplines are gradually integrated or discarded as needed [3]. This situation presents a significant challenge. Suppose a gradient-based optimisation identifies the local optimum in the initial design phase, taking into account disciplines that are too costly to consider at a later stage. When these outcomes are applied to a more intricate model, aiming to optimise additional design features and incorporating a different set of disciplines, potential superior designs for enhancing aircraft performance might be overlooked due to limited information and the "freezing" of the previous design optimisation. The sequential optimisation approach stems from the need for detailed models. Attempting to optimise all design variables and disciplines simultaneously (such as structural details, flight dynamics, etc.) is computationally impractical [4, 5]. Hence, multiple models are employed across distinct design phases.

Despite the utilisation of lamination parameters to condense the number of design variables, the Aeroelastic Tailoring problem remains high-dimensional, with a common design problem easily involving hundreds or thousands of design variables. These disciplines are accounted for by formulating G constraints for the optimisation problem (see Equation 1), arising from the multitude of analyses mentioned earlier. In such cases, the number of constraints is very large as well (when not relying on aggregation methods), further complicating the optimisation. This is why current design processes apply gradient-based local optimisation methods.

The present work addresses this problem by aiming to use multiple models of differing fidelity in a concurrent optimisation scheme to circumvent sequentiality and to ensure structural feasibility while keeping the computational process efficient and applicable. Thus, it is the objective to forge links between the single design stages to be able to exchange information and to obtain a better overall aircraft while decreasing the computational costs. Beyond that, some disciplines can not be com-

puted/ used in an optimisation on a high-fidelity model due to immense computational costs. Examples include control, aeroservoelasticity, and post-buckling. Thus, it is aspirable to be able to compute these constraints on the low-fidelity model while still other disciplines are updated by the high-fidelity one. The idea is that information from a low-fidelity (beam-based) and a high-fidelity (shell-based) model are used by employing a multi-fidelity approaches. The envisioned design cycle is depicted in Figure 1, comparing the past with the future. Currently, while gaining knowledge about the design

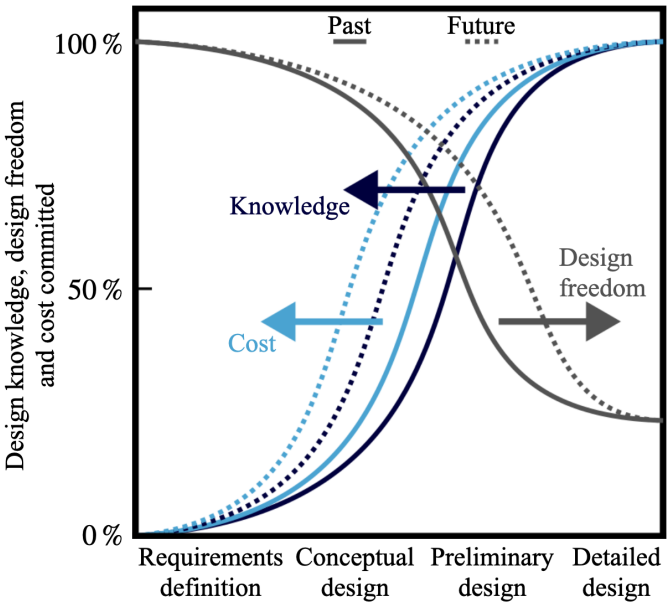

Figure 1 – Life-Cycle Design Stages (adopted from [3]).

with an ongoing design process, the design freedom is decreasing. For conventional designs the lack of knowledge within the initial design stages is compensated by the use of empirical knowledge. However, for novel aircraft configurations this empirical knowledge is lack and needs to be replaced by physics-based knowledge, further motivating the need for incorporating high-fidelity information as early as possible.

This work is structured as follows. First, multi-fidelity optimisation is introduced with a brief overview of the current state of the art. Next, the two models used in this work are presented, i.e. a beam-based and a shell-based model. The beam-based model will hereafter be referred to as the low-fidelity (LF) model, while the shell-based model will be referred to as the high-fidelity (HF) model. As stated earlier, especially in large aerospace projects, multiple models often exist and transit through the different design phases. These models are not necessarily derived from each other, as they may serve different purposes. A beam-based model does not require detailed input from a three-dimensional computer-aided design (CAD) model and can be used to predict the overall elastic and inertial behaviour of slender structures. However such beam-based models are limited in predicting failure modes that involve displacement fields that are more complex than what traditional beam formulations are capable of predicting, such as stress concentrations or distortions in cross-sections caused by local buckling or any other local phenomenon. Shell-based models are often developed once the CAD geometry is available and can be employed to produce a thin-walled representation of the 3D structure. Here, under the assumption of plane stresses, more detailed failure modes can be captured under the additional computational cost involved with orders of magnitude more degrees-of-freedom, when compared to the beam-based model.

In this study, we present an investigation comparing these two models and exploring their application in multi-fidelity design optimisation. After aligning the two models in terms of geometry, the results compare the responses to different structural and aeroelastic analyses to gain an understanding of their behaviour for later use in multi-fidelity optimisation. Following this, existing methodologies for multi-fidelity design optimisation are reviewed, and finally, some concluding remarks are presented.

2. Multi-Fidelity Multi-Disciplinary Design Optimisation

As introduced in the previous chapter, multi-fidelity design optimisation can significantly impact the design process of aerospace structures. In the following section, we briefly define the terminology of multi-fidelity and place it within the context of MDO (Multidisciplinary Design Optimisation). Generally speaking, the fundamental purpose of a model is to describe the relationship between an input and an output. In this work, we denote a model as a function $M : X \to Y$, mapping an input $x \in X \subseteq \mathrm{R}^d$ from the design space, where $d \in \mathrm{N}$, to an output $y \in Y \subseteq \mathrm{R}^g$, where $g \in \mathrm{N}$. Typically, this relationship is unveiled by solving a Partial or Ordinary Differential Equation. Often, multiple models describe the same physical system, each with different approximation accuracies and computational costs. High-fidelity (HF) models are typically defined as having high accuracy but being costly to evaluate. Conversely, low-fidelity (LF) models are less accurate but fast to evaluate [6]. As mentioned in Section 1, some quantities or constraints are too expensive to compute using the HF model but need to be respected concurrently during its optimisation.

Anyway, optimisation falls into the realm of many-query applications, demanding repeated model evaluations to find the optimal solution. Consequently, relying solely on HF models often becomes impractical due to their substantial computational demands. Conversely, performing optimisation solely based on the LF model accelerates computations. However, due to the lower and usually unknown approximation accuracy, this may lead to unsatisfactory results, such as infeasible or sub-optimal design points.

This motivates the development of multi-fidelity (MF) methods. These methods leverage the HF model to maintain accuracy while using the LF model, which captures the overall trend of the model output but may miss local details, to speed up computations. The goal is to obtain an optimisation result that is close to the outcome of a purely HF optimisation. This prompts the important question of how to manage these different models effectively.

A first and relatively simple approach is what the authors in [7] call fixed methods. Fixed methods have a predetermined pattern of LF and HF model evaluations. However, this approach requires a good correlation between the models in use. For problems with low correlation, convergence issues can arise. In other words, the optimiser may be distracted by a sudden evaluation of the HF model after several iterations of the LF model. Another multi-fidelity approach frequently used in the literature is to sequentially increase the level of fidelity, indicated by a predefined indicator, as proposed by [8]. The authors show that this approach can indeed decrease computational costs. These types of approaches are usually very economical since no further steps are needed to adjust the models. However, the aforementioned sequentiality remains present. To enhance the performance of multi-fidelity design optimisation, more sophisticated multi-fidelity approaches are required to be able to break out of the sequentiality.

[6] introduce three types of model management strategies aimed at balancing and correlating LF and HF model information to ensure accuracy while minimising computational cost. These three types are: *Adaptation* [9, 10], *Fusion* [11], and *Filtering* [12].

In addition to distinguishing between different strategies for managing MF models, LF models can be categorised based on their derivation. [6] define three different types of models: *Simplified models*, which use expertise from the HF model to derive an independent model with lower fidelity and differing computational costs; *Projection-based models*, which project the governing equations onto a low-dimensional subspace to yield a reduced model, exploiting the problem structure mathematically instead of relying on knowledge to derive a LF model; and *Data-fit models*, which are derived from input and output data, functioning as a black box.

Let in the following M_{lf} be the LF and M_{hf} the HF model. The mapping from the design space X to the outputs can be written as

$$\begin{aligned} M_{lf} &: X \to \{f_{lf}, \mathbf{c}_{lf}, \nabla f_{lf}, \nabla \mathbf{c}_{lf}\} \\ M_{hf} &: X \to \{f_{hf}, \mathbf{c}_{hf}, \nabla f_{hf}, \nabla \mathbf{c}_{hf}\} \end{aligned} \tag{2}$$

where $f(\mathbf{x}) \in \mathrm{R}, \nabla f(\mathbf{x}) \in \mathrm{R}^d$ is the objective function and its gradients and $\mathbf{c}(\mathbf{x}) \in \mathrm{R}^g, \nabla \mathbf{c}(\mathbf{x}) \in \mathrm{R}^{g \times d}$ the constraints and their gradients.

One of the earliest introductions to using multiple models within a concurrent local optimisation process can be found in the works of [13, 14]. The authors present the so-called Trust Region Model

Management (TRMM) strategy, which manages different numerical models based on the Trust Region approach. The standard Trust Region approach typically uses a quadratic approximation of the HF model, which is then optimised within the Trust Region. The authors in [13] generalise this idea by using the LF model instead, while occasionally evaluating the HF model to check progress.
Let $f_{hf}(\mathbf{x}i)$ be the objective function of the HF model and $flf(\mathbf{x}_i)$ the LF model at the i-th design point $\mathbf{x}_i$. If the LF model is constructed to satisfy first-order consistency, written as

$$\begin{aligned} f_{lf}^{(i)}(\mathbf{x}_i) &= f_{hf}^{(i)}(\mathbf{x}_i); \\ \tilde{\mathrm{N}} f_{lf}^{(i)}(\mathbf{x}_i) &= \tilde{\mathrm{N}} f_{hf}^{(i)}(\mathbf{x}_i) \end{aligned} \tag{3}$$

at every point $\mathbf{x}_i$ of the i-th iteration. Subsequently, the optimisation problem inside the trust region is solved to find the next update point. Depending on a comparison between the LF and HF models, the step size is adjusted. This approach, due to its first-order consistency, provably converges to an HF local optimum. Following this, the method is extended to maximise the use of the LF model to further reduce computational cost by presenting three nonlinear programming algorithms and demonstrating their applicability on an Euler analysis with varying degrees of mesh refinement [14].
Further works have built on this foundation. For instance, TRMM has been used with Bayesian model calibration to circumvent the need for HF derivatives [15] and has introduced a modified Trust Region method called Trust Region Filter SQP [16, 17]. To address differing numbers of design variables in LF and HF models, [18] apply a space mapping approach (first introduced by [19]) in combination with TRMM, mapping LF variables to HF ones. [20] then applied the space mapping approach to transonic airfoil aerodynamics.
[21] introduced the multi-fidelity Broyden–Fletcher–Goldfarb–Shanno (MF-BFGS) optimisation algorithm, based on a quasi-Newton solver. This method demonstrated faster convergence and reduced computational costs compared to TRMM when applied to an unconstrained problem [21]. The MF-BFGS method was then used to solve an aeroelastic design optimisation problem, optimising the range of an aircraft using seven design variables [22]. Subsequently, [23] investigated the scalability of the MF-BFGS approach.
However, all the aforementioned methods assume an unconstrained optimisation problem, where constraints are introduced via penalties added to the objective function to address this limitation. Penalty methods only find approximate solutions. However, when combined with gradient-based methods, they exhibit numerical issues [24]. Explicitly handling constraints, as done in Sequential Quadratic Programming (SQP) or within algorithms such as the Globally Converging Method of Moving Asymptotes (GCMMA) [25], has not yet been addressed in the literature.
As pointed out by [7], there is a lack of scalable methods that can handle high-dimensional inputs and outputs (large-scale constraints) without relying on aggregation methods, while efficiently leveraging gradient information. Engineering problems like the one introduced in Section 1 are often constrained in nature, and therefore it is expected that an accelerated convergence would be achieved by developing a novel multi-fidelity optimisation method that accounts for an explicit handling of constraints, which is the ultimate goal of this line of research. Furthermore, a multi-fidelity scheme that accounts for a subset of constraints, available only in one model, shall also be considered.

3. Aeroelastic Tailoring: Models

In line with prior discussions, multi-fidelity models usually stem from one another, as in *projection-based models*. However, as highlighted in Section 2. LF models also exist that are derived based on knowledge about the physical system, known as *simplified models*, without requiring complex 3D representations of the geometry before the models can be generated. This scenario commonly arises in large interdisciplinary aerospace projects, where diverse frameworks coexist and an increasing amount of information about the system becomes available throughout the various design phases. These models represent the same physical system but are derived from different formulations, typically aligning with the information available at various design phases. Each model has its own strengths and weaknesses. Efficiently handling and exploiting these diverse capabilities and the generated data, however, presents a significant challenge in today's aerospace industry. To leverage multi-fidelity optimisation with these independent models, this work aims to bridge this gap for the

use in Aeroelastic Tailoring.
Hereinafter, as previously introduced, two independent models are used. The HF model is a shell-based Nastran model, referred to as the Embraer Benchmark Wing, which serves as the ground truth. This model will be explained in more detail in Subsection 3.1
The LF model, however, is a beam-based model that is part of the low- delity design framework Proteus, developed by the Aerospace Structures and Materials department at Delft University of Technology, initiated with the work of [26], as described in Subsection 3.2
The overall optimisation problem is noted in Equation 1. The design variables are the lamination parameters (denoted with the superscript l) of each panel as well as their thickness, denoted with the superscript t.

$$x = \quad x_1^l; x_1^t; \dots; x_{n_p}^l; x_{n_p}^t \quad : \tag{4}$$

Note that in this work, the lamination parameters as part of the design variables are denoted by x_i^l, whereas in the composite community they are often referred to as $x1;2;3;4^{A;D}$. In the following subsections, both models are brie y explained.

3.1 EMBRAER Benchmark Wing: A High-Fidelity Aeroelastic Tailoring Model in Nastran

The HF model is the EMBRAER Benchmark Wing in NASTRAN. This model was developed to represent a generic aircraft, similar to the EMBRAER E-Jet family. It serves as a benchmark model for testing new analysis and design frameworks within an international Aeroelastic Tailoring consortium. In this model, the skin, spar, and rib covers are represented by shell elements, for which the material properties are speci ed separately via the matrices $A(x)$ and $D(x)$. The lamination parameters can be used with classical laminate theory to construct the following relationship:

$$\begin{matrix} N \\ M \end{matrix} = \begin{matrix} A(x) & B(x) \\ B(x) & D(x) \end{matrix} \begin{matrix} e^0 \\ k \end{matrix} : \tag{5}$$

The complete derivation of how to compute the matrix entries of $A(x)$ and $D(x)$ can be found in [27, 28]. These lamination parameters then de ne the material properties in the MAT2 card. Furthermore, additional parts such as pylon/engine and aileron assemblies are modelled using exible beam and shell elements to obtain a realistic representation. The aileron and other moving surfaces are likewise modelled via composite materials, similar to the main wingbox. Additionally, a realistic actuator system combined with a hinge mechanism is used. To ensure proper inertia properties of the system, the engine itself is modelled using rigid elements connected to the engine mount. Non-structural masses such as fuel tanks, high-lift devices, landing gear, accessories, and miscellaneous items are accounted for as lumped masses. This is crucial to maintain the correct dynamic behavior of the system. Similarly, the fuselage, payload, and tail are added as point masses and are used to obtain the proper trim condition.
The static and dynamic aerodynamic loads are computed using the Doublet Lattice Method (DLM) in NASTRAN. For further details, the reader is referred to [28]. As can be seen, the structure is much ner discretised, leading to an expectedly more accurate computation of eld quantities compared to the beam model, where information is condensed to the beam nodes. Additionally, more details are available, as shown in Figure 2. Of course, there is still room to further increase the delity of the model, for example, by using higher- delity aerodynamic solvers. However, it's important to note that this work focuses on preliminary aircraft design, where computationally ef cient models are still needed for use in a multi-query context. This model can then be used within NASTRAN's optimisation environment SOL200 to compute the corresponding constraints, arising for example from a buckling analysis with SOL105 or a utter analysis in SOL145.
This can be written as the input-output mapping, denoted as

$$\begin{aligned} M_{hf} : X \quad & ! f \quad f_{hf}; c_{hf}; \tilde{N} f_{hf}; \tilde{N} c_{hf} g \\ & f_{hf}(x) \; 2 \; R \quad c_{hf}(x) \; 2 \; R^g \quad \tilde{N} f_{hf}(x) \; 2 \; R^d \quad \tilde{N} c_{hf}(x) \; 2 \; R^{g \; d} \end{aligned} \tag{6}$$

Figure 2 – Aerodynamic planform and structural layout (from [28]).

Figure 3 – Planform measurements.

3.2 PROTEUS: A Low-Fidelity Aeroelastic Tailoring Framework

In general, the framework transforms a three-dimensional wingbox made of laminate panels into a nonlinear three-dimensional beam model. A panel in this setting can be, for instance, the upper and lower skin cover or the skin of the front and rear spar. Additionally, the wingbox can be discretised in the chord-wise direction with a user-de ned number to increase the design freedom. The airfoil shape, as can be observed in 5, is used here to size the wingbox. However, the panel method used to compute the aerodynamic loads is based on a at plate airfoil. The lamination parameters can be used with the classical laminate theory to construct the relationship de ned in Equation 5. This relationship encodes the dependency of the design variables with the stiffness of the system. A cross-section modeller [29], based on a variational approach, is used to obtain the element Timoshenko cross-sectional stiffness matrix $C \in \mathbb{R}^{6\ 6}$ by relating the strains e to the applied forces and moment s or $F_i; M_i$ respectively, as in

$$s = Ce\,! \quad F_1 \quad F_2 \quad F_3 \quad M_1 \quad M_2 \quad M_3 \ ^T = C \ e_{11} \quad e_{12} \quad e_{13} \quad k_1 \quad k_2 \quad k_3 \ ^T \tag{7}$$

with k_1 as the twist and $k_2; k_3$ the bending curvatures. Therewith, the properties of a 2D cross section are mapped onto the corresponding beam node, leading to a 6 6 element Timoshenko beam stiffness matrix. The corresponding beam strain energy in the continuous form can be derived like

$$U = \frac{l_0}{2} \int_0^1 e^T C e dx: \tag{8}$$

Figure 4 – Shell-based model in NASTRAN including aerodynamic panels.

Figure 5 – Beam representation of the wing structure. Diamonds represent the beam nodes. The beam nodes are chosen to be located at each rib, leading to 30 beam nodes and thus 29 beam elements.

By discretisation and introduction of the element degrees of freedom p 2 $R^{n_{dof}}$, the linear beam element stiffness matrix, at this point neglecting geometric and material nonlinearities, can be computed by

$$e = Bp \qquad K_{ij} = \frac{\P^2 U}{\P p_i \P p_j} \tag{9}$$

where the matrix B interpolates the nodal quantities to strains e. n_{dof} 2 N is the number of degrees of freedom (DoF), leading to 12 DoF per beam element, thus 6 DoF per node. The beam model is depicted in Figure 5, where diamonds represent the beam nodes.
Assuming a force vector f, the static solution is given as

$$K(p)p = f: \tag{10}$$

Note that geometrical nonlinearities are introduced by using the co-rotational framework of [30], decomposing large displacements and rotations into rigid body motion and small elastic deformations, ultimately leading to a dependence of the stiffness matrix K on the displacements p.
After formulating the nonlinear structure, the aerodynamic forces and moments are computed via the continuous-time state-space unsteady vortex lattice method (UVLM), derived in [31], and mapped onto the structure. Assuming low to moderate subsonic Mach numbers, very high Reynolds numbers as well as small angles of attack, the ow can be assumed to be incompressible, inviscid and