RESEARCH PAPER

Aerostructural topology optimization using high fidelity modeling

Pedro Gomes[1] · Rafael Palacios[1]

Abstract

We investigate the use of density-based topology optimization for the aeroelastic design of very flexible wings. This is achieved with a Reynolds-averaged Navier–Stokes finite volume solver, coupled to a geometrically nonlinear finite element structural solver, to simulate the large-displacement fluid-structure interaction. A gradient-based approach is used with derivatives obtained via a coupled adjoint solver based on algorithmic differentiation. In the example problem, the optimization uses strong coupling effects and the internal topology of the wing to allow mass reduction while maintaining the lift. We also propose a method to accelerate the convergence of the optimization to discrete topologies, which partially mitigates the computational expense of high-fidelity modeling approaches.

Keywords Topology optimization · Fluid-structure-interaction · MDO · Discrete adjoint · Multiphysics

1 Introduction

Topology optimization (TO) has been successfully applied to aircraft structures for more than one decade. So far, industrial-type applications have mainly used the decoupled approach where a particular component is designed in isolation and for fixed loads (Krog et al. 2004; Chin and Kennedy 2016; Zhu et al. 2016). Some researchers have started in parallel to explore TO based on fully-coupled simulations of fluid-structure interactions (FSI). The components targeted tend to be larger, such as wing boxes (James et al. 2014; Dunning et al. 2015; Kambampati et al. 2020) or ribs and spars (Maute et al. 2002; Maute and Allen 2004), and the optimizations have mainly focused on structural objectives (e.g. weight and compliance), while other variables, such as angle of attack or twist distribution, are used to maintain trim.

A shared characteristic of the previous examples is the use of medium-fidelity fluid models, namely vortex lattice methods or the Euler equations, and of linear elasticity for the structure. While they reduce the computational cost, lower-fidelity models can also result in unfeasible

✉ Rafael Palacios
  r.palacios@imperial.ac.uk

1   Imperial College, London SW7 2AZ, UK

designs. This is chiefly for two reasons. From an aerodynamic design point of view, TO applied to whole structure tends to produce surface irregularities as the wing deforms, as we have shown in Gomes and Palacios (2020). An inviscid fluid model does not penalize those solutions, but those irregularities may have a significant impact on viscous drag even under subsonic conditions. From a structural point of view, linear models artificially increase surface area as the wing deforms. A geometrically nonlinear structural model is then required to enforce inextensionality constraints. It also allows local buckling to be implicitly taken into account during the optimization process.

Other aerodynamically richer applications have been studied in the literature, showing that TO can be used to radically alter the aerostructural characteristics, e.g., to produce load alleviation (or augmentation), either passively or synergistically with an active system, or to obtain specific flutter characteristics (Maute and Reich 2006; Stanford and Ifju 2009; Stanford and Beran 2011; Stanford et al. 2013; Townsend et al. 2018; Stanford 2021). However, relatively small computational models have been used for these applications, since aerodynamic-driven problems create additional challenges for TO that ultimately increase the cost of the optimizations.

For density-based methods, a major challenge is that aerodynamic objectives do not necessarily benefit from a discrete material distribution (one without intermediate densities). In fact, even when lower fidelity models are used, the

distributed nature of aerodynamic loads makes the resulting topologies highly sensitive to the resolution of the grid, since it may be more efficient to support a pressure loading with a large number of small structures that a small number of large ones. Note that this mesh sensitivity aspect also affects level-set approaches (Kambampati et al. 2020). Taking viscous effects into consideration increases the challenge due to their close dependency with small irregularities on the wing surface, as we have mentioned above. Furthermore, as deformations increase, so does the influence of the internal structural topology on the aerodynamic quantities since, for example, the structure is able to adapt passively to different operating conditions. It is therefore important to understand the structural mechanisms that appear in optimizations to produce certain aerodynamic effects (for example, augmenting load), for they are likely to develop when multiple operating points are considered, and regardless of whether this is desired or not. Simultaneously, it is also relevant to propose methods to accelerate the convergence of this type of computationally expensive problem.

Here, these aspects are demonstrated by a flexible wing where, by optimizing its internal topology, the same aerodynamic performance can be maintained by an overall more flexible, hence lighter, structure. The work is carried out within the SU2 open-source suite, which provides a high-fidelity simulation environment that has been well described in the literature, both in terms of basic methods (Economon et al. 2016; Burghardt et al. 2020), and recent numerical improvements (Gomes et al. 2021) that have contributed to the feasibility of this work. Section 2 outlines the most relevant features of the models and their implementation. To alleviate the aforementioned discrete-topology challenges, a two-material topology is sought (similarly to Stanford (2021)), therefore the entire structure is solid, and guarantees that the outer layer is never left without support. Such a structure, built from a combination of flexible and stiff material, may also be easier to manufacture in certain conditions, than one with hollow regions (for example, if additive manufacturing techniques are used, auxiliary support structures are not needed since the two materials support each other). The focus of the numerical example in Sect. 3 is on the static characteristics of wings with large deformations, as a proof of concept of a new design strategy for wings built using additive manufacturing.

2 Methodology

A partitioned approach is used to solve the FSI problem. In particular, a Newton–Krylov method is used to solve the Reynolds-averaged Navier–Stokes (RANS) equations on the fluid side, Newton iterations are used for the nonlinear structural equations, and the fluid mesh is deformed based on a pseudo-elasticity problem (Sanchez et al. 2018; Gomes and Palacios 2020). The transfer of fluid loads and structural displacements is based on isoparametric interpolation. Gradient-based methods are used for numerical optimization since these are the only cost-efficient alternative for general large-scale optimization problems. The gradients of functions (objective and constraints) are obtained using a coupled discrete adjoint solver based on algorithmic differentiation (AD) (Albring et al. 2016; Burghardt et al. 2020). Similarly to the primal problem, a Krylov approach is used on the fluid side and a block Gauss–Seidel (BGS) method is used for coupling the fluid and structural problems. SU2 is used for primal and adjoint computations. The main methods and their implementation are described in the references above and are only succinctly described below. The emphasis of this section is on the key methodological improvements, in particular, the Krylov approaches, developed for this work.

2.1 Fluid problem

A RANS formulation is used to describe the fluid dynamics. A Newtonian fluid assumption is considered to compute the viscous stress tensor, and bulk viscosity effects are ignored. Turbulence is modeled as an increased viscosity using the "noft2" variant of the Spalart–Allmaras turbulence model with first-order convective fluxes (Spalart and Allmaras 1994). Pressure and temperature are related to the conservative variables via the ideal gas equation of state. Finally, a finite-volume discretization with median-dual grid is used, which results in a semi-discrete integral equation for each volume Ω_i of the form (Economon et al. 2016)

$$\int_{\Omega_i} \frac{\partial \boldsymbol{w}}{\partial t} d\Omega + \boldsymbol{R}_i(\boldsymbol{w}) = \boldsymbol{0}, \tag{1}$$

where $\boldsymbol{w}$ is the vector of conservative variables. In particular, convective fluxes are computed using Roe's scheme with 1% entropy correction, the second-order reconstruction of flow variables uses Green–Gauss gradients and Venkatakrishnan and Wang's limiter. The residual $\boldsymbol{R}_i$ is obtained by summing the discretized fluxes for all faces of the control volume and integrating the volumetric sources. To obtain a steady-state solution, equation (1) is marched implicitly in pseudo time (τ), that is, the new solution $\boldsymbol{w}^*$ is obtained by solving

$$\left(\frac{|\Omega_i|}{\Delta \tau_i} \delta_{ij} + \frac{\partial \tilde{\boldsymbol{R}}_i}{\partial \boldsymbol{w}} \right)(\boldsymbol{w}^* - \boldsymbol{w}) = -\boldsymbol{R}_i, \tag{2}$$

where the tilde indicates that the linearization of the residual is approximate. Equation (2) is solved using a backwards-Euler discretization in time. For second-order upwind schemes, the most common approximation is to consider only the influence of direct neighbors, which simplifies the

computation of the Jacobian matrix and reduces its density (especially for median-dual discretizations). However, this simplification can have the adverse effect of reducing the robustness of the solution method, which is highly relevant for numerical optimization to avoid divergence of intermediate designs. Moreover, adequate convergence of the primal problem has been shown to improve the convergence of the adjoint (Xu et al. 2015). Newton–Krylov (NK) methods can be implemented in a matrix-free manner by obtaining Jacobian-vector products using finite differences, the forward mode of AD, or the complex-step method. The accuracy of the Jacobian is then only limited by the differentiation strategy. For this work, a finite-difference-based NK method based on Wang et al. (2021) has been implemented for the flow equation, while turbulence is solved by a quasi-Newton method in a segregated manner. The linear preconditioner used for the NK approach is the ILU(0) factorization of the approximate Jacobian matrix. Currently, this is the bottleneck of the numerical implementation, namely with the maximum CFL number that can be used (≈ 100) it is not advantageous to use a fully coupled approach between flow and turbulence equations.

In what follows, we represent this fluid solution process by the fixed-point iteration

$$w = \mathcal{F}(w, u), \tag{3}$$

where u are the structural displacements, responsible for the deformation of the fluid mesh. Finally, and for the purposes of the fixed-point representation, the turbulence variables are considered to be part of w.

2.2 Solid mechanics

A finite-strain formulation is used in the solid domain, which is written in weak form as

$$\int_{\Omega_r} \delta \overline{\overline{E}} : \overline{\overline{S}} \, d\Omega - \int_{\Gamma_c} \delta u \cdot \lambda \, d\Gamma = 0 \tag{4}$$

where $\delta \overline{\overline{E}}$ is the variation of the Green–Lagrange strain tensor corresponding to δu, $\overline{\overline{S}}$ is the second Piola–Kirchhoff stress tensor and λ the external (fluid) tractions. Ω_r refers to the reference (undeformed) configuration and Γ_c to the current configuration on the surface. A library of materials has been implemented in SU2 and, for the analysis in this work, Neo-Hookean constitutive relations are considered (Bonet and Wood 2008). Equation (4) is linearized around the current configuration before being discretized with finite elements (Sanchez et al. 2016) (only linear elements are available in SU2). The solution is then iteratively found via the Newton–Raphson method, where the updates obtained with the tangent stiffness matrix are solved using the direct sparse solver in PaStiX (Hénon et al. 2002).

Noting that the tractions are computed based on the fluid variables, this process is also formulated as a fixed-point iteration, namely

$$u = \mathcal{S}(u, w). \tag{5}$$

The coupled problem is solved by a BGS approach by alternating between fluid (3) and structural (5) iterations. Both displacements and fluid tractions at the interface are interpolated using an isoparametric approach. A fixed relaxation factor (0.6–0.8) is applied to the transferred displacements to improve stability. We found this relaxation to be important during optimization iterations where the structure may become too flexible. However, designs that meet the target deformation constraints (described later) are not sensitive to the relaxation factor.

2.3 Coupled discrete adjoint sensitivities

Consider now a function of interest, J, which depends on the solution variables, $x = (w, u)$, and on the parameters of the problem, α, namely, undeformed mesh coordinates, local material properties, etc. From the solution of the FSI problem, it is $x = x(\alpha)$ and the total derivatives of J with respect to α can be obtained via the discrete adjoint method as (Albring et al. 2016; Burghardt et al. 2020)

$$d_\alpha J = J_\alpha + \lambda^T \mathcal{P}_\alpha, \tag{6}$$

where $\mathcal{P} = (\mathcal{F}, \mathcal{S})$ is the fixed-point iterator of the coupled problem, and subscripts denote partial derivatives. The adjoint variables $\lambda = (\overline{w}, \overline{u})$, where overbars denote the adjoints of the associated primal solution variables, are obtained by solving

$$\lambda^T = J_x + \lambda^T \mathcal{P}_x. \tag{7}$$

Similarly to the primal problem, the solution of the coupled adjoint problem can be obtained by a BGS method with inner iterations on the fluid and structural adjoint subproblems. Using the block structure of $\mathcal{P}_x$, we first rewrite (7) in partitioned form as

$$\begin{aligned} \overline{w} &= (J_w^T + \mathcal{S}_w^T \overline{u}) + \mathcal{F}_w^T \overline{w}, \\ \overline{u} &= (J_u^T + \mathcal{F}_u^T \overline{w}) + \mathcal{S}_u^T \overline{u}. \end{aligned} \tag{8}$$

Each subproblem in (8) is a fixed-point iteration with a contribution on the right-hand side due to the FSI coupling. For computational efficiency, those coupling terms are not updated while each discipline fixed point is converged. For example, the fluid contribution to the structural adjoint, $\mathcal{F}_u^T \overline{w}$, is only computed on the last inner iteration for $\overline{w}$. Noting that updating $\mathcal{F}_u$ involves a mesh deformation, this approach results in substantial computational savings overall.

This fixed-point formulation naturally fits the reverse mode of AD and, more importantly, allows a generic treatment of any type of solver (Burghardt et al. 2020) by differentiating its *entire* primal fixed-point iteration (the program, or code path, used to obtain a new solution iterate). There are however some drawbacks to this strategy. Effectively this solution is a right-preconditioned version of the more conventional left-preconditioned adjoint fixed-point approach (Giles et al. 2003). That means that the preconditioner is applied to the adjoint variables, instead of to the residual of the adjoint linear system. Note that the preconditioner is effectively the transposed (by AD) solution method of the recorded primal iteration, e.g. a quasi-Newton strategy. When this embedded preconditioner consists in the solution of a linear system by a Krylov solver, it is then necessary (in general) to converge such iterative processes to higher accuracy than in the primal solver (Gomes and Palacios 2021). Furthermore, it is well known that fixed-point approaches (either left- or right-preconditioned) may suffer from stability issues (Campobasso and Giles 2003; Xu et al. 2015; Kenway et al. 2019), even when the primal problem shows sufficient convergence. Krylov-based adjoints are commonly used to both accelerate convergence and improve stability. They can be seen as the dual operation of the NK methods for the primal (nonlinear) equations (Kenway et al. 2019). To cast (7) as a linear system, we simply write it as

$$(\mathcal{P}_x^{\mathrm{T}} - I)\lambda = -J_x^{\mathrm{T}}, \tag{9}$$

which can be solved by any matrix-free method, such as GMRES. However, if $\mathcal{P}_x$ is updated between iterations, it is not guaranteed that the Krylov solver will minimize the adjoint residuals. In Gomes and Palacios (2021) we have found that to solve (9) using GMRES, it is advantageous to use a fixed number of a Richardson iterations (20–30) with suitable relaxation (0.4–0.7) so that the residual is reduced between 1 and 2 orders of magnitude, as the smoother in the primal quasi-Newton process (2) recorded with AD (this is in contrast with the conventional primal strategy of using a Krylov method). Because this type of smoothing guarantees the linearity of $\mathcal{P}_x$, it makes the adjoint-GMRES approach highly effective for fluid problems (compared with the fixed-point approach). The structural adjoint problem can be solved by the fixed-point method without issues, its preconditioner (effectively the inverse of the tangent stiffness matrix) is very close to the inverse of the Jacobian $\mathcal{S}_u$, and thus around five iterations are typically sufficient.

2.4 Density-based topology optimization

A density-based approach with continuous variables is used in this work. It specifies a design density at discrete locations of the solid domain (the element centroids), with the local elasticity modulus being a function of it. Here, the modified SIMP (Simplified Isotropic Material with Penalization) formulation is used to relate the elasticity modulus with the design variable as described in Bendsøe (1989). For a two-phase material we have

$$E(\rho) = E_1 + (E_2 - E_1)\rho^p \quad 0 \le \rho \le 1, \tag{10}$$

where E_1 is the elasticity modulus of the more flexible material phase, and it is greater than zero even for solid-void problems to avoid a singular stiffness matrix. E_2 elasticity modulus of the stiffer material phase. To avoid checkerboarding in the solution, a discrete filtering operation (Bruns and Tortorelli 2001) is applied to the design density variables (ρ) exposed to the optimizer. This results in physical densities ($\tilde{\rho}$) considered for each finite element given by

$$\tilde{\rho}_i = \sum_{j \in \mathcal{N}(i)} w_{ij}\rho_j \quad / \sum_{j \in \mathcal{N}(i)} w_{ij}, \tag{11}$$

where $w_{ij} = (R - ||\mathbf{x}_i - \mathbf{x}_j||)\Omega_j$ and $\mathcal{N}$ is the set of elements within radius R of element i. This conical filter kernel is less numerically challenging than other strategies that introduce steep gradients as they approach discontinuous functions. However, it invariably results in regions of intermediate density.

The resulting problem is characterized by a large number of design variables and relatively few constraints (excluding simple bound constraints), which are imposed here using the exterior penalty method, by converting a conventional optimization problem, namely

$$\min_{\boldsymbol{\alpha}} f(\boldsymbol{\alpha})$$
$$\text{subject to} : g_i(\boldsymbol{\alpha}) = 0 \tag{12}$$
$$h_j(\boldsymbol{\alpha}) \le 0$$

into an unconstrained problem, that is

$$\min_{\boldsymbol{\alpha}} \quad f(\boldsymbol{\alpha}) + \sum_{i=1}^{N_g} a_i g_i(\boldsymbol{\alpha})^2 + \sum_{j=1}^{N_h} b_j h_j^+(\boldsymbol{\alpha})^2, \tag{13}$$

where $h^+ = \max(0, h)$. The penalized objective function can then be minimized using an unconstrained (but bounded) optimization method. The penalty parameters (a_i and b_j) need to be gradually increased (usually by multiplying the previous value by a fixed factor r) until a predetermined small constraint tolerance is met. This creates the need for outer iterations since updating the parameters within unconstrained (inner) iterations leads to bad approximations of the Hessian matrix. Although these outer iterations force an undesired reversion to steepest descent, they are also needed to update the TO parameters. However, in the example in this paper, this type of parameter update was not necessary. Before the constraint tolerance is met, loose convergence

criteria are used for the unconstrained optimizer (e.g. 40 inner iterations). The objective function is shifted and scaled by representative minimum value and range, respectively, while the constraints are shifted by their bounds and scaled by a reference value, which is chosen as the reciprocal of the bound unless otherwise specified. Doing so allows for a single constraint tolerance (≈ 0.01) to be used, and all penalty parameters to be initialized as $a_i^0, b_i^0 \in [1, 10]$ and updated with the same factor ($r \in [1.4, 4]$).

The unconstrained optimizer used in this work is the L-BFGS-B (Zhu et al. 1997), in the implementation of the SciPy library (Jones et al. 2001). This choice was made based on initial comparisons with the method of moving asymptotes of Svanberg (1987) and with the interior point optimizer IPOPT (Wächter and Biegler 2006). Despite some promising results in the literature (Rojas-Labanda and Stolpe 2015), IPOPT underperformed L-BFGS-B in our studies. Recently, other researchers have also found that current versions of IPOPT have difficulties with TO problems (Kennedy and Fu 2021). One computationally advantageous aspect of the exterior penalty approach, not explored in this work, is that it may reduce the cost of evaluating gradients. For example, it may be possible to differentiate the penalized function at the same cost of differentiating one of the constraints or objectives, and as a result, only one adjoint problem needs to be solved per primal simulation (e.g. operating point) instead of one per function.

3 Results

All the methods discussed above have been implemented and independently verified in the SU2 framework. This section reports on a first study on the feasibility of topology optimization at the wing level to achieve aerodynamic design objectives. This is an intermediate step towards full wing design, as the implementation is still limited to optimizing either structural or aerodynamic design variables (always on the coupled problem), but not both simultaneously as it will be needed in practice. In preparation for this work, we have obtained (Gomes et al. 2021) good numerical scalability for discretizations in the fluid and structural domains in the order of 10^7 and 10^6 nodes, respectively. This allows us to explore for the first time in the open literature coupled optimizations on wing topology including fluid viscosity. Resolving the boundary layer puts a physical constraint on the quality of the aerodynamic shapes. Otherwise, surface irregularities could easily appear after wing deformation due to non-conventional internal layout. However, the structural solver does not scale still at the level of Aage et al. (2017), which have used 10^9 nodes for a wing design (under fixed loads), and this sets a limit to the topological features that can be resolved. It should be also remarked that the

advantages of compliant wing design only become obvious for a multipoint optimization, as we have previously shown for 2-D problems (Gomes and Palacios 2020), since the solution to a single operating point is that of a wing that deforms to an optimal shape that can be obtained only via aerodynamic analysis. However, to limit the complexity in the explorations below, we have restricted ourselves to optimizations in a single operation point.

Taking into account the previous considerations, the following workflow has been established. (1) The starting point is a baseline design of an aeroelastic wing with a structure that fills the whole solid domain. A sufficiently stiff material is selected so that elastic displacements are small. (2) Shape optimization is then used (Sect. 3.1) to produce a minimum-drag design that will be used as reference. (3) A second wing is then designed for minimum weight using a less stiff material (Sect. 3.2). It is chosen to have the same undeformed aerodynamic shape as the first one, but, because of the softer material, the aerodynamic forces produce a negative elastic twist that reduces incidence and therefore lift. The classical solution to this problem would be to increase root angle of attack and then recompute the aerodynamic shape. Here we seek however to use TO to produce a load augmentation effect that recovers the desired lift as the structure is optimized. Viscous effects on the fluid still constrain the design to streamlined shapes with little or minimal separation, but drag is no longer part of the optimization since it would over-constrain the design (as we will discuss later). That would simultaneously need external shape design variables, which are not considered here.

3.1 Baseline near-rigid design

A baseline design is obtained by optimizing the outer shape of a solid cantilever wing with uniform material properties for minimum drag, with constraints on the lift, pitching moment, and minimum thickness (over multiple spanwise sections). The initial geometry is chosen using simple design rules and it is shown in Fig. 1. It results from linear interpolation between two symmetric 4-digit
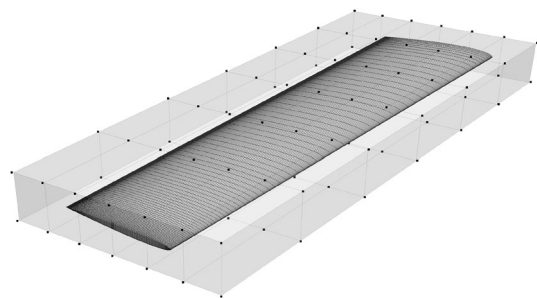

Fig. 1 Initial wing geometry and FFD box for shape optimization

NACA profiles. The root profile has 0.25 m chord and 9% thickness, whereas the tip has 0.175 m chord and 7.2% thickness. The wingspan is 1 m, the 25% chord-line is swept back 5°, and a linear twist distribution of −3° is used. The free-form deformation (FFD) box shown in Fig. 1 is used to parameterize the shape. Out of the 98 control points visible in the figure, 97 are allowed to move in the vertical direction (±30 mm), whereas the bottom right point of the root section (right side of the page) is fixed to prevent vertical translation in the optimization. Chord and translation (chord-wise displacement) are defined for each of the 7 spanwise sections of the FFD box by controlling the horizontal coordinates of the 14 control points in each section. The chord is changed by moving the points away from the centre-line of the FFD box, proportionally to their distance to the line. The tip section is not allowed to translate, this helps to maintain the quality of the mesh. The bounds for these two variables are ±60 mm at the root and decrease linearly to ±20 mm at the tip. There are 110 variables in total, and note that these are able to change the effective angle of attack by pitching the wing, despite the flow direction not being changed by the optimizer. A single operating point is considered, namely Mach 0.6 at sea-level and 273.15 K. The corresponding Reynolds number at the wing root is 3.7 M. The fluid and structural grid are composed mostly of hexahedra and have 4.1 million nodes and 220,000 nodes, respectively. The radius of the farfield boundary is 25 span lengths, Fig. 2 shows a detail of the structural grid near the wingtip. The solid
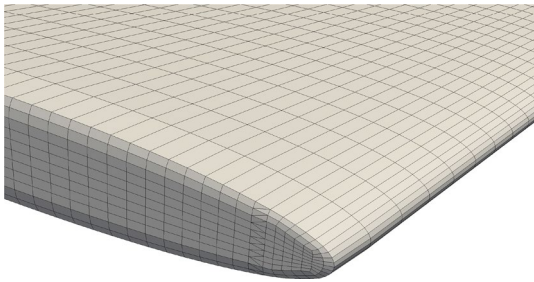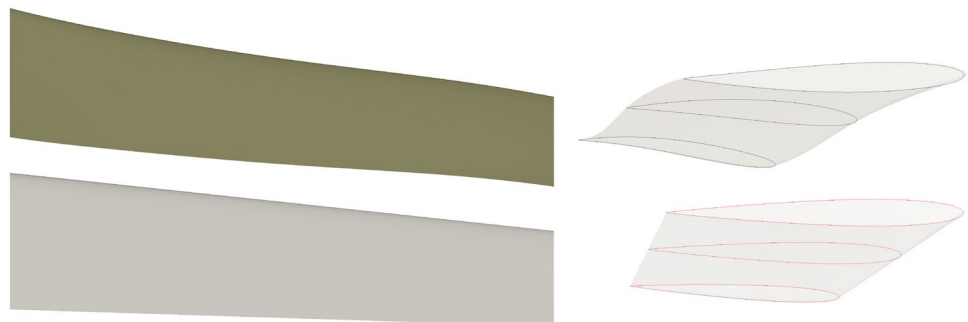

Fig. 2 Structural mesh near the wingtip

material has Poisson's ratio of 0.35 and Young's modulus of 75 GPa, which results in deformations comparable to the wing thickness.

At 4° angle of attack, the initial geometry has coefficients of lift, drag, and pitching moment of 0.295, 0.0103, and −0.0429 respectively. To obtain a baseline for TO, drag (C_D) is minimized with lower bounds of 0.29 on lift (C_L) and −0.06 on pitching moment (C_M) coefficients. A geometric constraint was also included to prevent the maximum thickness of the thinnest spanwise section (t_{min}) from decreasing. In summary,

$$\min_{\boldsymbol{\alpha}} C_D(\boldsymbol{\alpha})$$
$$\text{subject to}: C_L(\boldsymbol{\alpha}) > 0.29$$
$$C_M(\boldsymbol{\alpha}) > -0.06 \qquad (14)$$
$$t_{min}(\boldsymbol{\alpha}) > 0.072$$

The SLSQP implementation from SciPy is used to solve the optimization problem. All constraints and variables are scaled by their bounds and the objective (drag) by its initial value.

Figure 3 shows a comparison of the planforms and spanwise sections at the root, mid-span, and tip of the initial and optimized geometries. All constraints are active on the optimized design and the drag coefficient is reduced by nearly 20% to be 0.00829. The design achieves this by reducing the loading on tip sections and increasing it in the first half of the span to produce a more elliptical distribution. The sections towards the tip are also significantly cambered, which requires the root sections to operate at higher incidence to counteract the effect of camber on pitching moment, as shown by the pressure coefficient contours in Fig. 4. However, note that despite the increase in incidence near the root, on the optimized design those sections produce less lift due to a reduction in chord. Similarly, the sections near the tip produce more lift due to camber, despite the lower incidence. This is shown by the lift distributions in Fig. 5. As expected when minimizing drag at subsonic conditions, the result is a more elliptical lift distribution. The fluid domain is finally remeshed for

Fig. 3 Comparison of initial (bottom) and optimized (top) geometry for the baseline wing

Fig. 4 Pressure contours on initial (bottom) and optimized (top) designs for the baseline wing
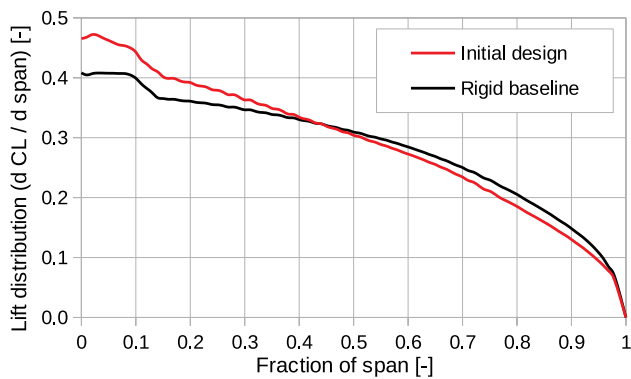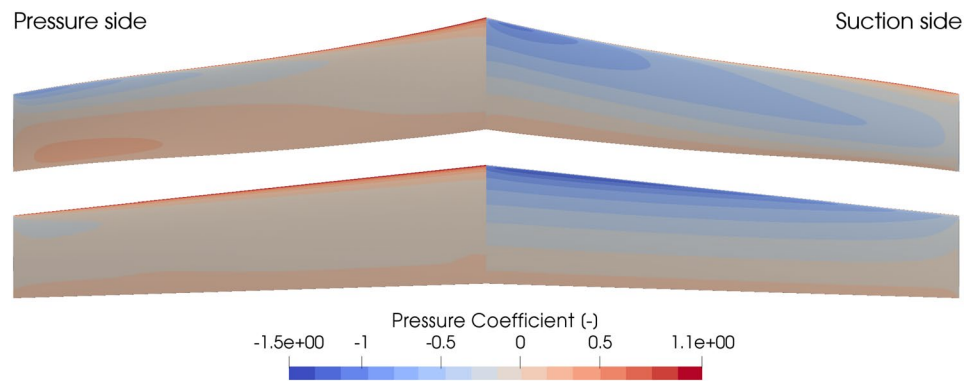
Fig. 5 Spanwise lift distribution on initial (light color) and optimized (dark color) designs

the optimized baseline design, before proceeding to optimize the internal topology of the wing.

3.2 Topology optimization

As it was discussed in the introduction, a TO problem is considered that divides the domain between a soft and a stiff material, without voids. It is not relevant for the purpose of this study if such materials exist and they are chosen instead such that the tip deformation is comparable to the chord, when using 50% fraction of stiff material. This results in a elasticity modulus of 1.5 and 30 GPa for the soft and stiff materials, respectively. The structural elements are stretched (see Fig. 2), because using a smaller size in the spanwise direction would result in an impractical number of elements for the current setup. This presents a challenge for density filtering strategies, since setting the filter radius based on the longest element size would span the entire thickness of the wing. To avoid this, the neighborhood search ($\mathcal{N}$ in (11)) was constrained to consider only neighbors of neighbors, on a 2-D grid this would limit the filter to the equivalent of a 13-point stencil. The filter radius was then set as 1.5 times the largest element size (in the spanwise direction), and the

linear weight function (11) was used. Finally, a fixed SIMP exponent of 3 is used throughout the investigation.

We commence with a solid wing with $E = 22.3$ GPa, for which the lift coefficient drops to 0.24. The first objective is then to recover the target value of 0.29, using the minimum amount of stiff material. Note that a density of zero in (10) corresponds to the softer material and a density of one to the stiffer material, and therefore minimizing the fraction of the latter is equivalent (in formulation) to minimizing mass in a classical "solid-void" problem. Note also that it is not possible to meet the constraint by stiffening the structure, since the stiffer material in the TO is still 60% more flexible than the one used for the baseline design. For a feasible design, some structural constraint needs to be introduced and, although maximum stress would be most appropriate, in our current implementation we seek to limit compliance. Moreover, given the challenges associated with aerodynamic-driven TO that have been described above, we have divided the optimization process into several steps, also as a way to make the trade-offs between goals more evident. These steps are discussed in the next subsections and the starting point for each step is the result from the previous unless noted otherwise. While this may lead to sub-optimal results, such incremental strategies are necessary to understand the challenges in such a large design space.

3.2.1 Step 1: recovering lift

The first step is to recover the target lift, which also serves as evidence for the feasibility of the optimization problem. To that end, the exterior penalty algorithm (13) is used with a single equality constraint with a fixed penalty factor of 32. Note that the algorithm would start by recovering lift even if other goals were included, since the initial design is infeasible and thus lift would be weighted heavily by the exterior penalty formulation. Figure 6 shows the surface view of the resulting topology after the 11 L-BFGS-B iterations required to meet the constraint, starting with a uniform fraction of 0.9. By removing stiff material close to the leading edge, the incidence on spanwise sections close to the root
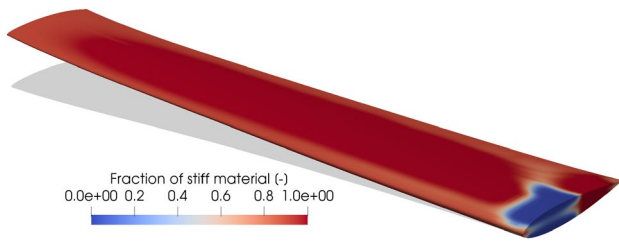
Fig. 6 Material distribution to meet lift constraint

negates the pitch-down moment of the tip sections, as shown by the distribution of aeroelastic twist (i.e. relative to the jig shape, and thus the variation in AoA due to the deformation) in Fig. 7. This is also necessary to compensate the smaller projected area that results from larger displacements. This geometrically-nonlinear effect is captured in our structural model but would be neglected by linear theories.

3.2.2　Step 2: determining a suitable bound for the fraction of stiff material

In the previous step, there was no incentive to use a minimal amount of stiff material. That is introduced in this step alongside an upper bound on compliance (W) of 75 J (the compliance after the initial step is 44.8 J). The optimization problem is

$$\min_{\rho} \frac{1}{V_\Omega} \int_\Omega \tilde{\rho}(\rho)\,\mathrm{d}\Omega$$
$$\text{subject to} : C_L(\rho) > 0.29 \tag{15}$$
$$W(\rho) < 75\,J$$

Although it is known that minimizing mass for given compliance may give rise to structures that are close to a stability limit, this formulation helps to determine a suitable target for the fraction of stiff material, which can then be constrained instead of outright minimized. The penalty factors for the
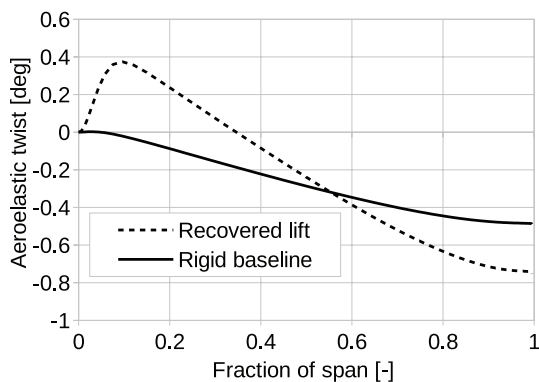
two constraints are fixed at 8 and the objective (fraction of stiff material) is not scaled since its initial value, which has resulted from the previous step, is already close to one. Figure 8 shows the resulting topology after 55 L-BFGS-B iterations. Mass fractions below 0.5 have been made transparent to reveal the internal structure, which is mostly composed of the softer material. In the figure, the spanwise slices are taken from 0.25 to 0.55 span in increments of 0.075. The main mechanism to maintain lift remains the same and the stiffer material was removed from less strained areas, e.g. close to the wingtip. The volume fraction of stiff material has been reduced in this optimization step to 0.33, at which point the compliance constraint is active. However, the material distribution is not yet discrete, mainly for two reasons: first, due to the relatively small number of iterations (for a topology problem); and second, due to the lift constraint, which requires material to be removed from a high strain region (the root). Note that the lift constraint is responsible for the main load applied to the structure (and thus in part for its compliance) and this contradiction between both constraints in (15) makes the problem difficult to solve.

3.2.3　Step 3: lift-constrained compliance minimization

To avoid the previous conflict between constraints, the roles of compliance and stiff material fraction are now switched, i.e. compliance is minimized under the 0.29 lift coefficient constraint and a maximum stiff fraction of 0.4. The problem then becomes

$$\min_{\rho} W(\rho)$$
$$\text{subject to} : C_L(\rho) > 0.29 \tag{16}$$
$$\frac{1}{V_\Omega} \int_\Omega \tilde{\rho}(\rho)\,\mathrm{d}\Omega < 0.4$$

Note that if this maximum value of stiff material fraction were known a priori, e.g. from experience, step 2 could have
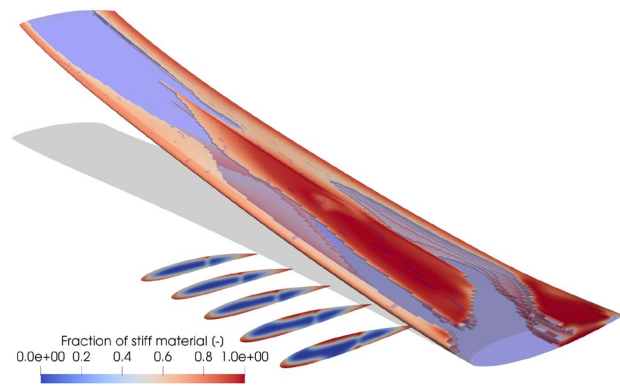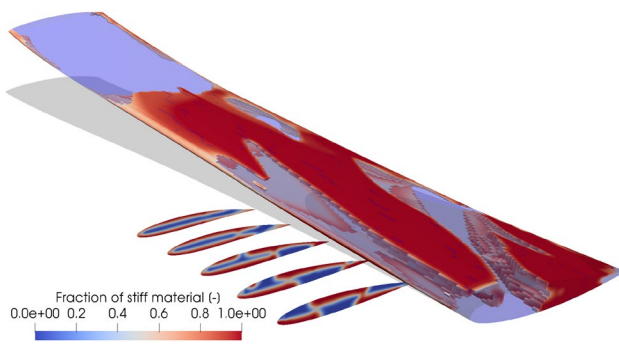

Fig. 7 Spanwise distribution of twist to meet lift constraint

Fig. 8 Material distribution after minimizing stiff fraction

Fig. 9 Material distribution after minimizing compliance

Fig. 10 Material distribution after minimizing stiff fraction for a lower compliance constraint

been avoided. The initial penalty factors were again set to 8 but increased by 50% (up to 64) every 30 L-BFGS-B iterations for constraints that did not meet a 0.5% tolerance. Note that constraints are scaled by their bounds and the objective by its initial value. Figure 9 shows the resulting topology after 225 iterations, which has a compliance of 38.1 J, lift coefficient of 0.286, and stiff fraction of 0.403. The material distribution is more discrete than before, note the smaller regions of intermediate density. However, only 50% of the variables are at their bounds (0 or 1), whereas in a typical topology optimization problem all variables converge to either the lower or upper bound.

3.2.4 Step 4: improving the discreteness of the material distribution

Ideally, a TO problem should benefit from a discrete solution due to the way it is posed and parameterized. However, in this problem, the requirement on lift makes this difficult to achieve, since the structure is both responsible for creating the load and resisting it. It is worth showing that posing the problem as compliance (and lift) -constrained mass minimization is not beneficial. To that end, the roles of stiff fraction and compliance were switched (recovering equation (15)). However, the penalty factors are now fixed at 64 and the bound on compliance lowered to 40 J aiming to direct the optimization towards the current local optimum (i.e. the result of step 3) which was also used as the starting point.

Figure 10 shows the resulting topology after 180 iterations. The stiff material fraction is reduced to 0.336 with lift coefficient of 0.285 and compliance of 40.2 J. However, it is clear that conventional formulations, based on compliance and material fraction, do not produce a discrete result for this problem. Along with reducing the size of areas with intermediate material fraction (note the larger gaps from Figs. 9 and 10), i.e. moving the interfaces, the optimization has also removed stiff material from areas that were previously solely stiff (note the change near the leading edge at 60 to 70% span on the same two figures). A
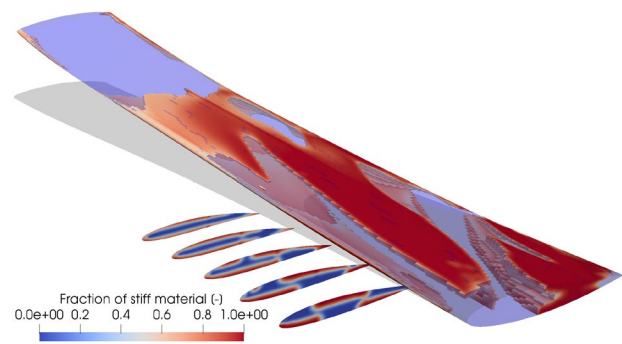
plausible explanation for this is that, to support a distributed load, it is better to distribute the stiff material than to concentrate it to produce stiff regions.

In the TO literature, authors have used explicit penalization of non-discreteness metrics for problems that do not converge naturally to discrete solutions (see e.g. Chen and Wu (1998)). It is known that such strategies need to be managed with care (e.g. ramped), to avoid fast convergence to locally optimal solutions before any topological features begin to develop. However, ramping strategies have the downside of increasing the computational cost of the optimization, which is already significant for this problem due to high-fidelity modeling. Moreover, even for topology problems that naturally converge to discrete solutions, most of the objective function reduction takes place in the first iterations, where the main features of the topology develop. The remainder of the iterations serve mostly to refine those features, and in general, do not contribute substantially towards improving the objective function. For example, note how in this problem the lift constraint can be met by a wide range of topologies (from mostly soft to mostly stiff). In other words, once the main topological features are defined, the problems becomes relatively insensitive to the exact location of the interface between material phases, provided the structure is not close to buckling. Based on these observations it is tempting to propose an early termination of the optimization, followed by a post-processing operation to force the discreteness of the result (for example considering all values above a certain threshold to be completely stiff and vice versa). Though this may be viable for more linear problems, in the presence of strong FSI it is likely that some constraints will no longer be respected after such a naive post-processing.

Therefore, to address the challenges of high computational cost and of obtaining discrete topologies, the solution proposed here is as follows. After obtaining a feasible solution (e.g. either the current one or also the one in step 3), the discreteness of the solution is improved by explicitly targeting (minimizing) a non-discreteness metric (d),

namely a version of that used in Sigmund (2007) adapted to unstructured meshes,

$$d = \frac{1}{V_\Omega} \int_\Omega 4\tilde{\rho}(1 - \tilde{\rho}) \, d\Omega, \quad (17)$$

where $\tilde{\rho}$ are the filtered densities (material fractions in this case) and N the number of design variables. The reduction of computational cost is achieved by simulating only the structure under fixed fluid loads (obtained from the last FSI simulation). To ensure those loads are appropriate throughout the entire process, a deformation target is introduced for the surface (Γ) nodes, that is,

$$\epsilon_\Gamma = \frac{||\boldsymbol{u}_\Gamma - \boldsymbol{u}_\Gamma^*||^2}{N_\Gamma}. \quad (18)$$

This error metric is minimized (as a secondary objective) to obtain a small value (10^{-7} in this case). These are the only two objectives used in this step. Despite the similar formulation, this step has a very different role than the inverse-design approach of previous work (Gomes and Palacios 2020). Here, the initial topology is nearly discrete, and thus only requires some incremental refinement, while the inverse approach in our previous work does not use initialization (i.e. it starts from uniform density) and therefore may produce substantially different material layouts.

Stiff fraction minimization for a compliance of 40 J has resulted in a discreteness metric of 0.3. This is reduced to 0.2 over the course of 1000 (inexpensive) iterations, after which the error metric is 4×10^{-8} and the stiff fraction 0.339. To ensure that the final solution would be as discrete as the type of filter allows, the weight of the non-discreteness metric was progressively increased up to 500. In the end, 75% of the variables were within 2% of their bounds. Furthermore, evaluating (17) for the design variables (i.e. unfiltered fractions) yields 0.008. This indicates that to lower the value of d further would require a different type of density filter (e.g. morphology-based filters, as in Sigmund (2007)). As expected, the lift constraint is not respected by this solution, for which the value is 0.272. To restore it to the target ($C_L = 0.29$) the discretized solution was used as starting condition for a coupled (FSI) optimization, where non-discreteness is minimized, i.e.

$$\min_{\boldsymbol{\rho}} d(\boldsymbol{\rho})$$

$$\text{subject to}: C_L(\boldsymbol{\rho}) > 0.29 \quad (19)$$

$$\frac{1}{V_\Omega} \int_\Omega \tilde{\rho}(\boldsymbol{\rho}) \, d\Omega < 0.35$$

After 25 iterations both this metric and the stiff fraction were maintained, the lift was restored but compliance increased to 42.3 J. Compliance was not constrained or included in

the objectives since this would recover (in part) the formulation that did not result in a discrete material distribution. This shows the key challenge associated with aerodynamic objectives, namely that even introducing compliance-minimization aspects into the problem (either as objectives or contraints) may not contribute to obtaining a discrete solution due to the dependence of the fluid loads on the structural response. This is shown here by the trade-off between compliance and a more discrete material distribution. The resulting material distribution is shown in Figs. 11 and 12.

3.2.5 Discussion

The optimization process described above can be summarized in the following steps:

1. Determining the feasibility of the aerodynamic goals, e.g. meeting lift ($O(10)$ iterations).
2. Approximately solving the optimization for suitable structural targets ($O(100)$ iterations). This may include determining those targets via different formulations of the problem (as done here).
3. Accelerating convergence to a discrete solution considering only the structure under fixed loads and target shape (equivalent to $O(10)$ iterations).
4. Correcting the aerodynamic goals for the discrete solution ($O(10)$ iterations).

Most of the optimization effort takes place in step 2, while steps 3 and 4 effectively are post-processing operations. Moreover, note that step 1 would not be necessary if aerodynamic constraints are guaranteed by other variables (e.g. shape or angle of attack). Table 1 summarizes the evolution of the key performance values over the different optimization steps.

The drag coefficient for the design that only meets the lift constraint (step 1) is just over 10% higher than in the baseline design. While some of that increase is due to the
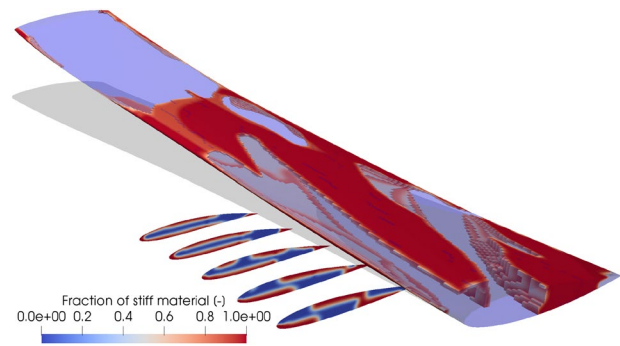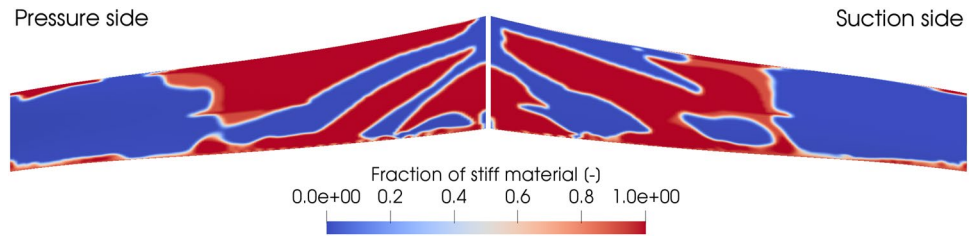

**Fig. 11** Material distribution after improving discreteness (3-D view)

**Fig. 12** Material distribution after improving discreteness (surface view)

Pressure side   Suction side

Fraction of stiff material (-)
0.0e+00   0.2   0.4   0.6   0.8   1.0e+00

Table 1 Summary of key performance values during the optimization steps

| Step	C_D	C_L	Comp. (J)	Stiff frac.
1. Recover lift	0.00917	0.290	44.8	0.938
2. Min. compl.	0.0109	0.286	38.1	0.403
3. Min. stiff frac.	0.0113	0.285	40.2	0.336
4. Final design	0.0115	0.290	42.3	0.338

mesh distortion, most of it comes from the larger deformation, which reduces the planform area and thus results in a higher wing loading than the one the wing shape has been optimized for in Sect. 3.1. This is expected to become worse with larger deformations at smaller fractions of stiff material. It is worth noting that single-point optimum shapes for flexible wings tend to result in flat deformed shapes (i.e. straight line from root to tip), which should be expected since this maximizes the projected area for a given span. For the designs with smaller stiff fraction (from step 2 onwards), the drag coefficient increases further, and in the final discrete solution (that meets the lift constraint more strictly) it is almost 40% higher that in the baseline design. This is a significant increase that is not entirely justified by larger deformations. Recall that the lift-generation mechanism used by the TO solutions is to allow the mid-span to pitch up to both reduce and compensate for the effects of higher pitch-down moments near the tip. In addition, reducing the stiff fraction with limited compliance requires material to be removed from low strain regions. Note in particular how past 60% span the design uses almost no stiff material. Consequently, this allows those areas to deform more, that is, the incidence is reduced due to the pitch-down moments. This reduces the lift generated by those sections of the wing, in turn requiring the inboard sections to produce even more lift to compensate (drag increases as the spanwise distribution of lift moves away from the optimized baseline). This effect is visible in Figs. 13, 14 and 15, where the deformation and lift distribution of the final design are compared with the one from step 1, and it explains the vestigial areas of stiff material near the tip of the wing, whose placement at the leading and trailing edges maximizes torsional stiffness, thereby preventing more negative incidence.
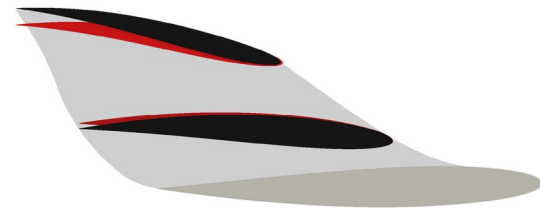
Fig. 13 Deformation at mid-span and tip, of the final design (dark colors, reduced by 20%), and the design obtained to meet the lift constraint (light colors)
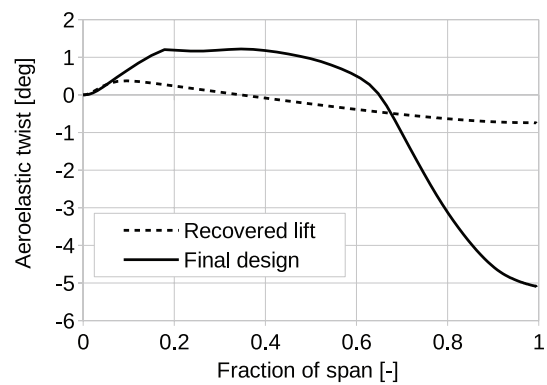
Fig. 14 Spanwise distribution of twist for the final design compared with lift-feasible design
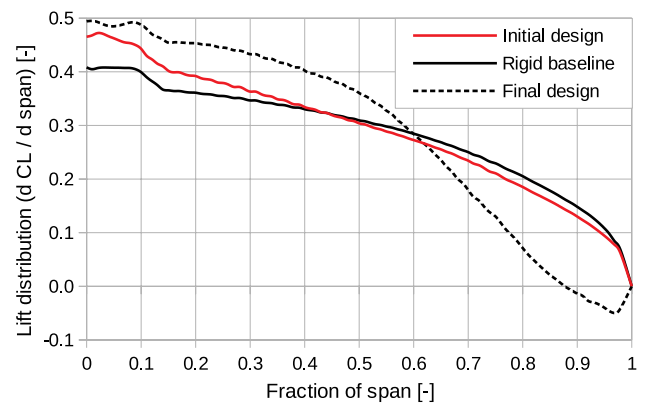
Fig. 15 Spanwise lift distribution on initial (light color), shape-optimized (dark color), and final (dashed line) designs

Note finally that these reinforced areas (discussed above) are not connected to the larger areas of stiff material, in effect, this reduces the coupling between the two (note the rapid change of twist after 0.6 span in Fig. 14). Moreover, although the final design deforms more in torsion and bending, the corresponding moments are lower. These interruptions of the main load paths (i.e. stiff material portions) are necessary for the aerodynamic goals but raise structural concerns. In particular, they cause the soft material to be more strained than its stiff counterpart, since it is compressed and sheared between floating regions of stiff material that are not connected to the root. To illustrate this, Fig. 16 shows the von Mises stresses divided by the local elasticity modulus, which makes apparent the need for stress constraints. Those were not considered here but will be explored in further iterations of this research.

4 Final remarks

This work aimed to explore the merits and challenges of applying topology optimization to aeroelastic wing design, that is, a technique that is well established for purely structural applications to problems where the structure is largely responsible for the aerodynamic performance. This was done by restricting the example problem to a single operating point and only structural design variables. While for a full assessment of TO as a wing design tool, many other aspects would have to be considered (shape optimization, multiple operating points, off-design performance, composite materials, manufacturability, cost, etc.), this use of TO appears to have merit, to the extent that the topological features can be traced back to aerodynamic effects. Furthermore, such features develop in a similar number of iterations as they would for purely structural problems. However, as a consequence of the limited means the optimization was given to realize its goals, the aerodynamic performance of the final design was reduced. Better performance could be expected, for example, from simply allowing the angle of attack to change, thereby allowing the inboard sections to operate at higher incidence without having to compromise the stiffness of the structure so severely by removing stiff material near the root. However, in so doing, the (desired) load augmentation characteristic

of the problem would be lost. Ideal performance, on the other hand, can only be expected from a simultaneous optimization of shape and topology, which has so far proved challenging with high-fidelity simulations. In particular, the exterior penalty approach is much less suited for shape optimization than SLSQP, especially if the design space is augmented with three orders of magnitude more design variables. Conversely, SLSQP is not adequate for large-scale optimization problems due to its dense approximation of the Hessian matrix.

A major challenge in this type of topology optimization, is the higher computational cost of using high-fidelity models. Approximately 10 times more (coupled) iterations are required than for aerodynamic shape optimization. The greatest challenge, however, is that aerodynamic design objectives do not benefit from an optimally stiff structure, and thus the optimization will not converge to a discrete solution unless that goal is explicitly introduced. Of course, level-set topology optimization methods (which were not explored in this work) guarantee that the solution will be discrete. However, a further contributor to the discreteness challenge is that loads are distributed over a large surface, whose smoothness is important for aerodynamic performance. Since it is more efficient to maintain that smoothness with a large number of small supports, than with a few highly rigid connections, the mesh size required to resolve those features may render 3-D applications impractical. In that respect, density-based methods may have some advantages due to their connection with homogenization theory. A recent research direction is the combination of level set and density-based methods to avoid the shortcomings of each method. This is achieved by obtaining a sharp interface with the level set approach, while using the density-based strategy to control feature sizes and seed hole creation (Andreasen et al. 2020; Barrera et al. 2020).

The proposed solutions address both the high computational cost and the discrete-solution challenges. Comparing level set with density-based methods for the type of problem studied here should be the subject of future work. For that, it will be also important to study the interaction between shape and topology optimization, whether by proposing efficient sequential approaches or by devising an optimization method that can consider both sets of variables simultaneously.

Fig. 16 Von Mises stress normalized by local elasticity modulus
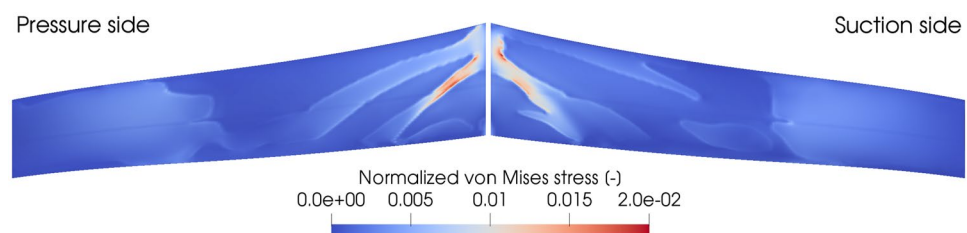
Declarations

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

Replication of results The input files for the results in this work can be found in https://doi.org/10.5281/zenodo.5188519.

References

Aage N, Andreassen E, Lazarov BS, Sigmund O (2017) Giga-voxel computational morphogenesis for structural design. Nature 550(7674):84–86

Albring T A, Sagebaum M, Gauger N R (2016) Efficient aerodynamic design using the discrete adjoint method in SU2. In: 17th AIAA/ISSMO multidisciplinary analysis and optimization conference

Andreasen CS, Elingaard MO, Aage N (2020) Level set topology and shape optimization by density methods using cut elements with length scale control. Struct Multidisc Optim 62(2):685–707. https://doi.org/10.1007/s00158-020-02527-1

Barrera J, Geiss M, Maute K (2020) Hole seeding in level set topology optimization via density fields. Struct Multidisc Optim 61:1319–1343. https://doi.org/10.1007/s00158-019-02480-8

Bendsøe MP (1989) Optimal shape design as a material distribution problem. Struct Optim 1(4):193–202. https://doi.org/10.1007/BF01650949

Bonet J, Wood RD (2008) Nonlinear continuum mechanics for finite element analysis, 2nd edn. Cambridge University Press, Cambridge

Bruns TE, Tortorelli DA (2001) Topology optimization of non-linear elastic structures and compliant mechanisms. Comput Methods Appl Mech Eng 190(26–27):3443–3459

Burghardt O, Gauger NR, Gomes P, Palacios R, Kattmann T, Economon TD (2022) Discrete adjoint methodology for general multiphysics problems. Struct Multidisc Optim 65(28):1–14. https://doi.org/10.1007/s00158-021-03117-5

Campobasso MS, Giles MB (2003) Effects of flow instabilities on the linear analysis of turbomachinery aeroelasticity. J Propul Power 19(2):250–259

Chen TY, Wu SC (1998) Multiobjective optimal topology design of structures. Comput Mech 21(6):483–492

Chin TW, Kennedy G (2016) Large-scale compliance-minimization and buckling topology optimization of the undeformed common research model wing. In: 57th AIAA/ASCE/AHS/ASC structures, structural dynamics, and materials conference

Dunning PD, Stanford BK, Kim HA (2015) Coupled aerostructural topology optimization using a level set method for 3D aircraft wings. Struct Multidisc Optim 51(5):1113–1132

Economon TD, Palacios F, Copeland SR, Lukaczyk TW, Alonso JJ (2016) SU2: an open-source suite for multiphysics simulation and design. AIAA J 54(3):828–846

Giles MB, Duta MC, Muller JD, Pierce NA (2003) Algorithm developments for discrete adjoint methods. AIAA J 41(2):198–205

Gomes P, Palacios R (2020) Aerodynamic-driven topology optimization of compliant airfoils. Struct Multidisc Optim 62(4):2117–2130

Gomes P, Palacios R (2021) Pitfalls of discrete adjoint fixed-points based on algorithmic differentiation. AIAA J 1–6

Gomes P, Economon TD, Palacios R (2021) Sustainable high-performance optimizations in SU2. In: AIAA Scitech 2021 Forum

Hénon P, Ramet P, Roman J (2002) PaStiX: a high-performance parallel direct solver for sparse symmetric definite systems. Parallel Comput 28(2):301–321

James KA, Kennedy GJ, Martins JR (2014) Concurrent aerostructural topology optimization of a wing box. Comput Struct 134:1–17

Jones E, Oliphant T, Peterson P (2001) SciPy: open source scientific tools for Python. http://www.scipy.org/

Kambampati S, Townsend S, Kim HA (2020) Coupled aerostructural level set topology optimization of aircraft wing boxes. AIAA J 58(8):3614–3624

Kennedy G, Fu Y (2021) Topology optimization benchmark problems for assessing the performance of optimization algorithms. In: AIAA Scitech 2021 Forum

Kenway GK, Mader CA, He P, Martins JR (2019) Effective adjoint approaches for computational fluid dynamics. Prog Aerosp Sci 11:100542

Krog L, Tucker A, Kemp M, Boyd R (2004) Topology optimisation of aircraft wing box ribs. In: 10th AIAA/ISSMO multidisciplinary analysis and optimization conference arXiv:1011.1669v3

Maute K, Allen M (2004) Conceptual design of aeroelastic structures by topology optimization. Struct Multidisc Optim 27(1–2):27–42

Maute KK, Reich GW (2006) Integrated multidisciplinary topology optimization approach to adaptive wing design. J Aircr 43(1):253–263

Maute K, Nikbay M, Farhat C (2002) Conceptual layout of aeroelastic wing structures by topology optimization. In: 43rd AIAA/ASME/ASCE/AHS/ASC structures, structural dynamics, and materials conference, Denver, Colorado

Rojas-Labanda S, Stolpe M (2015) Benchmarking optimization solvers for structural topology optimization. Struct Multidisc Optim 52(3):527–547

Sanchez R, Palacios R, Economon TD, Kline HL, Alonso JJ, Palacios F (2016) Towards a fluid-structure interaction solver for problems with large deformations within the open-source SU2 suite. In: 57th AIAA/ASCE/AHS/ASC structures, structural dynamics, and materials conference

Sanchez R, Albring T, Palacios R, Gauger NR, Economon TD, Alonso JJ (2018) Coupled adjoint-based sensitivities in large-displacement fluid-structure interaction using algorithmic differentiation. Int J Numer Methods Eng 113(7):1081–1107

Sigmund O (2007) Morphology-based black and white filters for topology optimization. Struct Multidisc Optim 33(4–5):401–424. https://doi.org/10.1007/s00158-006-0087-x

Spalart PR, Allmaras SR (1994) One-equation turbulence model for aerodynamic flows. Recherche Aerospatiale 1:5–21

Stanford B (2021) Topology optimization of low-speed aeroelastic wind tunnel models. In: AIAA Scitech 2021 Forum

Stanford B, Beran P (2011) Optimal structural topology of a platelike wing for subsonic aeroelastic stability. J Aircr 48(4):1193–1203

Stanford B, Ifju P (2009) Aeroelastic topology optimization of membrane structures for micro air vehicles. Struct Multidisc Optim 38(3):301–316

Stanford B, Beran P, Kobayashi M (2013) Simultaneous topology optimization of membrane wings and their compliant flapping mechanisms. AIAA J 51(6):1431–1441

Svanberg K (1987) The method of moving asymptotes-a new method for structural optimization. Int J Numer Methods Eng 24(2):359–373

Townsend S, Picelli R, Stanford B, Kim HA (2018) Structural optimization of platelike aircraft wings under flutter and divergence constraints. AIAA J 56(8):3307–3319

Wächter A, Biegler LT (2006) On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming. Math Program 106(1):25–57

Wang L, Diskin B, Nielsen E J, Liu Y (2021) Improvements in iterative convergence of FUN3D solutions. In: AIAA Scitech 2021 Forum

Xu S, Radford D, Meyer M, Müller JD (2015) Stabilisation of discrete steady adjoint solvers. J Comput Phys 299:175–195

Zhu C, Byrd RH, Lu P, Nocedal J (1997) Algorithm 778: L-BFGS-B: Fortran subroutines for large-scale bound-constrained optimization. ACM Tran Math Softw 23(4):550–560

Zhu JH, Zhang WH, Xia L (2016) Topology optimization in aircraft and aerospace structures design. Arch Comput Methods Eng 23(4):595–622