

Article

Aerostructural Design Exploration of a Wing in Transonic Flow [†]

Nicolas P. Bons ^{*} and Joaquim R. R. A. Martins 

Department of Aerospace Engineering, University of Michigan, Ann Arbor, MI 48109, USA; jrram@umich.edu

^{*} Correspondence: nbons@umich.edu

[†] This paper is an extended version of our paper published in 2020 AIAA Scitech Conference.

Received: 6 July 2020; Accepted: 4 August 2020; Published: 14 August 2020



Abstract: Multidisciplinary design optimization (MDO) has been previously applied to aerostructural wing design problems with great success. Most previous applications involve fine-tuning a well-designed aircraft wing. In this work, we broaden the scope of the optimization problem by exploring the design space of aerostructural wing design optimization. We start with a rectangular wing and optimize the aerodynamic shape and the sizing of the internal structure to achieve minimum fuel burn on a transonic cruise mission. We use a multi-level optimization procedure to decrease computational cost by 40%. We demonstrate that the optimization can transform the rectangular wing into a swept, tapered wing typical of a transonic aircraft. The optimizer converges to the same wing shape when starting from a different initial design. Additionally, we use a separation constraint at a low-speed, high-lift condition to improve the off-design performance of the optimized wing. The separation constraint results in a substantially different wing design with better low-speed performance and only a slight decrease in cruise performance.

Keywords: aerostructural optimization; design exploration; MDO; wing design

1. Introduction

1.1. Background

The aerodynamic behavior of a wing is tightly coupled to its structural response, and vice versa. Before the advent of modern multidisciplinary analysis, predicting the aeroelastic response for a given wing was limited. This limitation led to the development of the traditional wing design process, in which successive designs are passed iteratively between aerodynamics and structures engineering groups. In his seminal book on aircraft design, released posthumously in 1978, Kuchemann [1] writes:

This should be one of the aims for the future: we want an integrated aerodynamic and structural analysis of the dynamics of the flying vehicle as one deformable body, and to use that for design purposes.

The pursuit of this ideal gave birth to the field of multidisciplinary analysis and optimization (MDAO), wherein integrated aerostructural analysis and design framework is now a reality. Initially, simplified models were used out of necessity because of computational constraints. Haftka [2] combined a lifting line model with a simple finite-element model to perform one of the earliest aerostructural optimizations. Low-fidelity models continue to be used to facilitate analysis and optimization. Chittick and Martins [3] used a panel method and a single tubular spar to demonstrate aerostructural optimization. Jansen et al. [4] used an aerodynamic panel method and an equivalent beam structural model to enable exploration of the nonplanar wing design space using a gradient-free optimizer (which would require too many function evaluations to use with higher-fidelity models).

More recently, Jasa et al. [5] developed an open-source aerostructural model (OpenAeroStruct) in OpenMDAO [6] that uses a vortex-lattice method for aerodynamics and a beam finite-element model for the structure.

The objective of multidisciplinary design optimization (MDO) is to optimize the design parameters of multiple disciplines simultaneously, rather than sequentially. For MDO to fully replace the traditional wing design process, it must be capable of handling hundreds of design variables. Additionally, high-fidelity computational models are necessary to capture the influence of the design variables on the wing's performance. For a given optimizer, the number of iterations required to reach a solution increases with the number of design variables. In general, the cost of each optimization iteration increases with the fidelity of the computational analysis being used. Thus, for high-fidelity wing design, it is advantageous to use an optimizer that can handle many design variables while keeping the number of iterations low. Gradient-based optimizers require far fewer iterations to reach a solution than gradient-free methods [7].

Gradient-based optimizers are faster than gradient-free methods because they use derivatives to determine their path through the design space. However, derivative computations can be expensive. Additionally, the accuracy of the computed derivatives is critical to the success of the optimization. A naïve implementation of derivative computations might use the finite-difference method, which does not generally yield accurate gradients and has a computational cost that is proportional to the number of design variables [8]. There are more sophisticated approaches to gradient calculation, such as the complex-step approximation [9] and the adjoint method [10,11]. The adjoint method computes derivatives with the same level of accuracy as the primal solver and has a computational cost that is independent of the number of design variables.

The coupled-adjoint for aerostructural systems enables high-fidelity gradient-based optimization of realistic wing designs [12]. One of the first high-fidelity aerostructural optimizations was conducted by Martins et al. [13], who optimized the wing shape and wingbox sizing of a supersonic business jet using Euler computational fluid dynamics (CFD) and a finite-element model. Since then, there have been various other efforts using CFD-based aerostructural optimization with both Euler [14,15] and Reynolds-averaged Navier–Stokes (RANS) models [16–20]. Although more accurate fluid flow models are possible with large-eddy and direct-numerical simulations, the computational cost of such methods renders them prohibitive for wing design optimization with the current technology. Furthermore, RANS is accurate enough for drag minimization at cruise flight conditions [21,22]. Thus, RANS coupled with finite-element structural analysis represents the state-of-the-art for aerostructural wing design optimization.

The convergence of a gradient-based optimizer is determined by the Karush–Kuhn–Tucker (KKT) conditions, which ensure that the constraints are satisfied and the objective cannot be locally improved at the final solution. However, despite the rigor of the KKT conditions—or perhaps because of it—gradient-based optimizers are only guaranteed to converge to a single, local minimum. In optimization problems with multiple local minima, a gradient-based optimizer converges to only one solution, which may not be the global optimum. Multiple research efforts have shown that aerodynamic shape optimization (ASO) problems with airfoil shape and wing twist are unimodal [7,21,23,24]. However, the appearance of spurious, multiple local minima in these types of problems is possible when the convergence criterion of both the functions and derivatives is not sufficiently stringent, as discussed by Yu et al. [7]. These spurious local minima were also exposed by Koo and Zingg [24] in a follow-up to a previous paper [25].

When planform variables are added to the design problem, multiple local minima do appear in the design space [23,26–28]. However, it is crucial to distinguish mathematically rigorous local minima and physically significant local minima. For example, Chernukhin and Zingg [23] found many local minima in a benchmark design problem with chord, dihedral, sweep, and span variables, in addition to airfoil shape and twist. However, they used Euler CFD in the optimization, thereby creating a nonphysical design space with local minima that might not exist in reality. Streuber and Zingg [27]

and Bons et al. [26] approached the same benchmark problem using RANS CFD, but still found multiple local minima. However, by carefully studying the influence of each of the design variables, Bons et al. [26] discovered a physically legitimate reason for multimodality in the chord distribution. By adding a constraint to enforce a monotonically decreasing chord distribution, some multimodality was eliminated from the problem, and more practical designs were obtained. Other common examples of legitimate, multiple local minima include upward and downward winglets, and forward and aft swept wings.

An understanding of multimodality in the wing design space enables designers to pose optimization problems to facilitate design space traversal from the starting design to the global optimum. Many of the optimization problems in the literature involve small changes between the baseline and optimized designs. These refining optimization problems do not demonstrate the optimizer's ability to traverse the design space, as would be required to find a solution in the design of an unconventional aircraft. Instead, researchers often resort to randomly perturbing the initial design [21,27] or starting from a blank slate design [23,26]. These design space exploration studies bolster confidence in the suitability of gradient-based optimizers for ASO problems. However, they have not been replicated for aerostructural wing design problems. In the same way that RANS optimization results supersede Euler-optimized designs, the introduction of structures into the design problem creates an entirely new—and more realistic—design space to study. In the current work, we apply similar methods to a more realistic transonic wing design problem with consideration of structures. Adding a wingbox structure allows the optimization to find the proper trade-off between weight and drag as it varies the planform and nonplanarity of the wing.

Single-point optimizations are prone to exhibit poor off-design performance. One of the most common solutions for this problem is to set an objective function that is a weighted average of the performance at multiple design points. The set of design conditions included in the objective is referred to as a *multipoint* stencil. Thus, even though the optimization problems solved in Section 3.2 includes cruise, maneuver, and buffet analysis points, we designate them as single-point designs because the objective function was only based on a single design point. Using a multipoint objective improves the average performance across the stencil at the expense of the nominal design point. However, it can result in intermittent performance, wherein the design functions optimally at the specified design conditions but poorly in the intervals. Drela [29] reported this phenomenon in a set of airfoil optimization studies and showed that increasing the number of points in the stencil helped curb this tendency. In wing ASO, Lyu et al. [21] obtained a more robust design using a 5-point stencil than with a single-point optimization. The multipoint design had a weak shock across the stencil, whereas the single-point design had completely eliminated the shock at the nominal design point. They also found that the multipoint design had a larger leading-edge radius than the single-point design. Kenway and Martins [30] compared different multipoint stencils of varying size and composition and found a good compromise between robustness and computational expense with a carefully chosen 5-point stencil. Various other efforts have performed multipoint ASO successfully [31–35]. Multipoint optimization has also been demonstrated in aerostructural wing design [18,36,37].

Although there have been extensive comparisons between single-point and multipoint designs with ASO, the same cannot be said for aerostructural wing optimization. Additionally, most of the past efforts on multipoint design have focused on robust cruise performance without considering the impact of design changes at low-speed, high-lift conditions. Preserving low-speed, high-lift performance in a wing optimization problem is notoriously difficult because of the complications that arise from modeling and parameterizing high-lift devices. The difficulties associated with high-lift devices can be avoided by considering clean wing performance at low-speed, high-lift conditions. To this end, Wakayama and Kroo [38] and Ning and Kroo [39] have shown that constraining $C_{L,max}$ using critical section theory results in a more practical planform design. In airfoil optimization, Buckley et al. [40] added a constraint on $C_{l,max}$ into the multipoint objective function to meet safety requirements at a low-speed condition. Rather than constraining $C_{L,max}$, Khosravi and Zingg [14] included climb drag

in the multipoint objective function to encourage improvement in that regime. The whole issue is often skirted by simply imposing limitations on the geometric parametrization to prevent changes that would adversely affect high-lift performance (e.g., minimum leading-edge thickness). This work introduces a new approach to preserving low-speed, high-lift performance while minimizing cruise fuel burn.

1.2. Problem Description

The Common Research Model (CRM) was designed as a benchmark for the verification and validation of CFD solvers across industry and academia [41]. Subsequently, the CRM wing was adopted as the test case for a series of benchmark aerodynamic shape optimization problems by the American Institute of Aeronautics and Astronautics (AIAA) Aerodynamic Design Optimization Discussion Group (ADODG) [21]. More recently, Brooks et al. [18] reverse engineered the CRM to create an undeflected (jig) shape of the CRM as a benchmark for aerostructural analysis and optimization called the undeflected Common Research Model (uCRM). In this work, we start with a rectangular wing and solve the same optimization problem as the uCRM. We created a rectangular wing with the same reference area and aspect ratio as the uCRM-9, which we call the “plank”. The cross-section of the wing is the RAE 2822 airfoil with a trailing edge thickness of 5 mm. The planform and cross-section of the wing are shown in Figure 1 and the initial geometric properties are listed in Table 1.

Table 1. Rectangular wing specifications.

Property	Value	Units
Reference area	383.12	m ²
Half-span	29.38	m
Aspect ratio	9.01	
Mean aerodynamic chord	6.52	m
Sweep	0	degrees




Figure 1. Rectangular wing definition.

The objective of the optimization problem is to minimize fuel burn for a specified mission length. The optimizer is free to change the shape of the wing and the sizing of the wingbox to achieve minimum fuel burn. Fuel burn is calculated as:

$$\begin{aligned} W_3 &= W_{\text{fixed}} + W_{\text{wing}} + W_{\text{payload}} + W_{\text{reserve}} \\ W_2 &= W_3 \exp\left(\frac{R c_t C_D}{V C_L}\right) \\ W_{\text{fuel}} &= W_2 - W_3. \end{aligned} \quad (1)$$

The parameters W_{wing} , C_L , and C_D are subject to change during the optimization; the fixed properties are defined in Table 2 and V is based on the Mach number and altitude. The fixed mass, W_{fixed} , includes all aircraft components except the wing, payload, and fuel.

Table 2. Aircraft specifications.

Property	Description	Value	Units
W_{fixed}	Fixed mass	100,000	kg
W_{payload}	Payload	34,000	kg
W_{reserve}	Reserve fuel	15,000	kg
R	Mission range	7250	nm
c_t	Thrust-specific fuel consumption	0.53	h^{-1}

2. Methods and Tools

2.1. Computational Framework

We use the MDO of aircraft configurations with high fidelity (MACH) framework to solve the optimization problem introduced above. The component hierarchy and process flow for aerostructural optimization in the MACH framework are shown in the extended design structure matrix (XDMS) [42] diagram in Figure 2. At each iteration, the optimizer changes the design, the MDA solver converges the aerostructural system, and functions of interest are returned to the optimizer. In this work, we use the optimizer SNOPT v7.7 [43]. The geometry is parametrized with a free-form deformation (FFD) scheme [44] implemented in pyGeo [45]. The FFD parametrization applies to both the aerodynamic and structural meshes, so that the wingbox is always consistent with the outer mold line (OML). Changes to the OML are propagated from the aerodynamic surface nodes to the volume mesh using the inverse-distance mesh-warping algorithm in IDWarp. We obtain the solution of the aerostructural system with a Gauss–Seidel iterative scheme. ADflow [11,46] is used to obtain a RANS solution of the flow with the Spalart–Allmaras turbulence model. The Toolkit for the Analysis of Composite Structures (TACS) [47] is used to solve for the displacement of the structure under the aerodynamic loads. A Krylov method is used to solve the coupled adjoint of the multidisciplinary system. For this study, we consider the solutions of both the MDA and coupled adjoint sufficiently converged when the l^2 norm of the residual has decreased by 10^{-5} . The structural node displacements and aerodynamic surface loads are transferred between the aerodynamic and structural meshes using a rigid link load and displacement transfer scheme first introduced by Brown [48] and subsequently implemented in MACH [16].

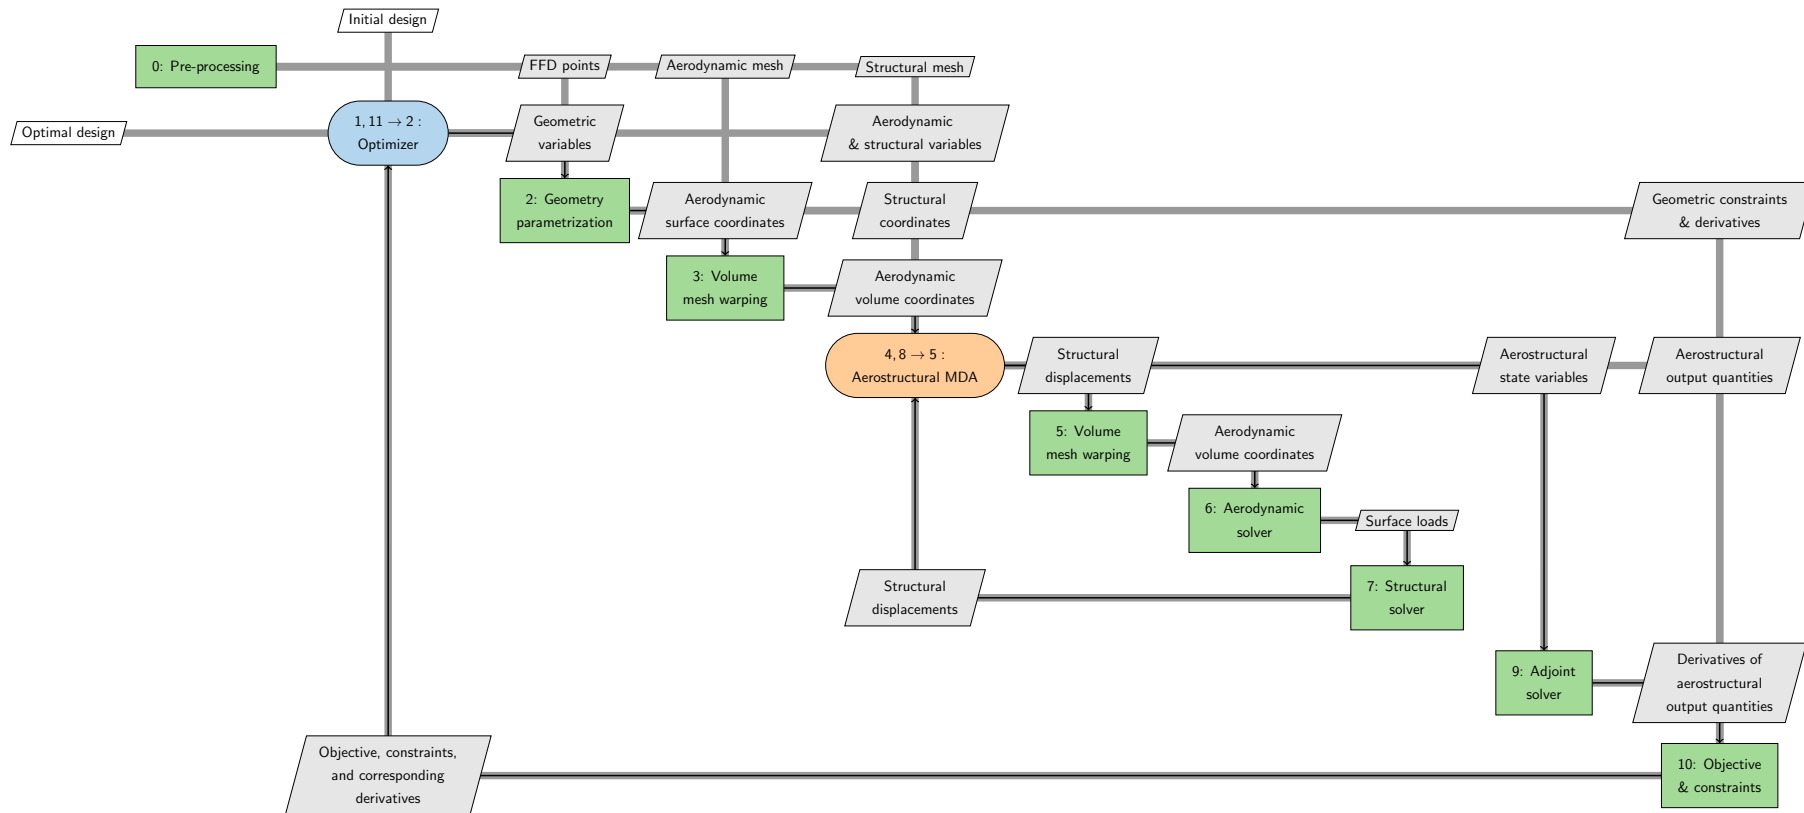


Figure 2. XDSM diagram of aerostructural optimization with MACH.

2.2. Preprocessing

2.2.1. CFD Meshing

We created two series of three meshes, for a total of six, shown in Figure 3. The finest mesh (A0) was created first and then coarsened by a factor of $\sqrt{2}$ to get the starting mesh for the B-series (B0). The level 1 and 2 surface meshes in each series were coarsened by factors of 2 and 4, respectively, from the level 0 mesh. All surface meshes were extruded to a far-field distance of $100 \times \text{MAC} = 652 \text{ m}$ with an initial off-wall spacing of $5.74 \times 10^{-6} \text{ m}$ using pyHyp. For these meshes, x is the streamwise direction, z is the lift direction, and a symmetry boundary condition is placed at $y = 0$. The dimensions of the meshes are listed in Table 3.

The purpose of making multiple CFD meshes is two-fold. First, it allows us to ensure mesh independence in a mesh convergence study, which is presented and discussed in Section 3.2. Second, multiple meshes are used for each optimization problem to reduce the overall computational cost. This multi-level optimization strategy is described and demonstrated in Section 3.1.

Table 3. Mesh dimensions.

Label	N_{edge}	N_{chord}	N_{span}	$N_{\text{off-wall}}$	N_{total}	Max y^+
B2	2	22	33	32	52,096	2.15
A2	3	32	48	44	152,064	1.42
B1	4	44	66	64	416,768	1.19
A1	6	64	96	88	1,216,512	1.12
B0	8	88	132	128	3,334,144	1.15
A0	12	128	192	176	9,732,096	1.15

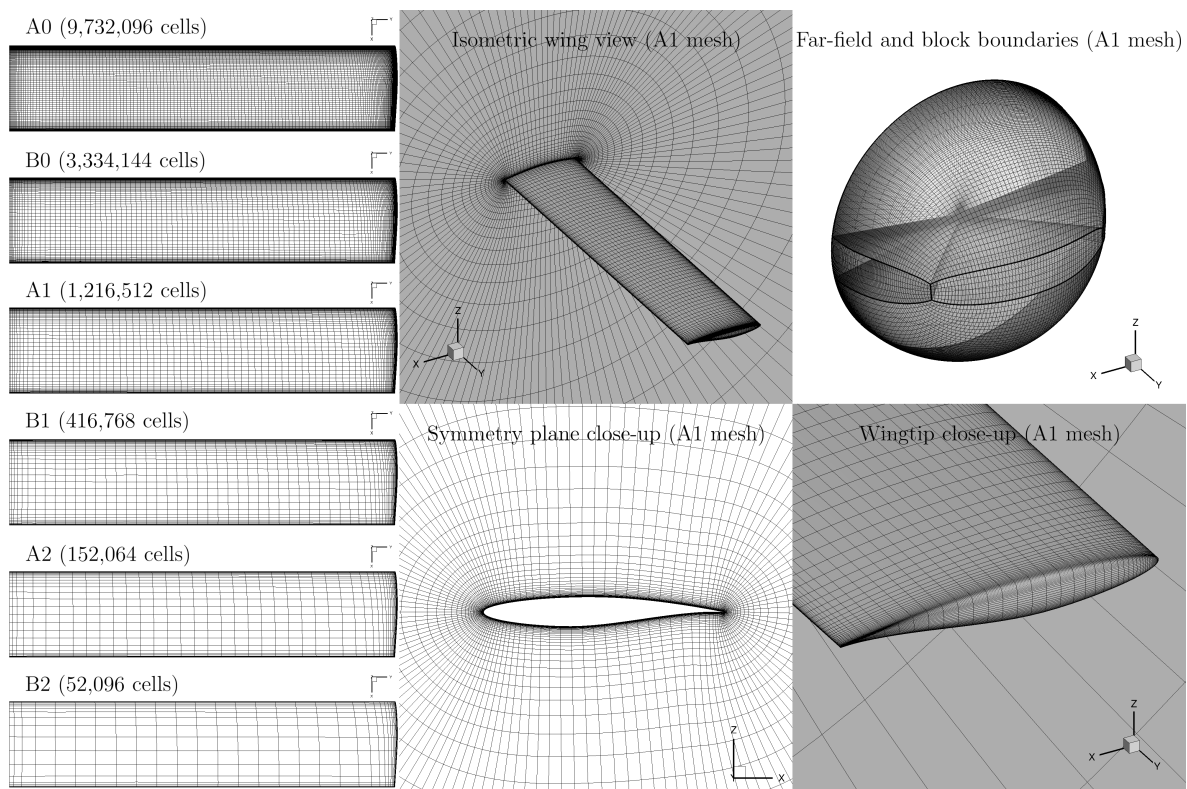


Figure 3. CFD meshes of baseline wing.

2.2.2. FEA Meshing

We patterned the structural mesh after the uCRM-9 wingbox. The fore and aft spars are located at 10% and 65% chord, respectively. The wingbox has 46 ribs, extending from the symmetry plane to

the wingtip. Figure 4 shows the placement of the wingbox within the wing and the discretization of the wingbox. The structural mesh consists of 16,672 quadrilateral, 2nd-order, MITC, shell elements.

Each component of the wingbox model is divided into seven panels in the chordwise direction and 45 panels in the spanwise direction (for a total of 630, 322, and 90 panels for the skins, ribs, and spars, respectively). Each of these panels consists of multiple shell elements. The panels of the skins and spars are modeled using the smeared stiffness approach described by Kennedy and Martins [49]. Each panel can have its own variables to control panel thickness, stiffener thickness, stiffener height, and stiffener pitch, as shown in Figure 5. The panels can also be grouped into design variable groups so that they share the same values for these parameters.




Figure 4. The FEA mesh has 16,672 2nd-order shell elements.




Figure 5. Smeared stiffness panel components.

2.2.3. Geometric Parametrization

As previously mentioned, the geometry is parametrized using an FFD volume. We use a coarse FFD for the optimizations on the B2 and A2 meshes and a fine FFD on the finer meshes. The coarse FFD has five control points distributed with cosine spacing along the chord and nine control points distributed evenly along the span, making a total of 90 FFD points (including top and bottom). The fine FFD has double the number of chordwise control points, for a total of 180. The coarse FFD is needed on the coarse CFD meshes to maintain an appropriate ratio of CFD points to FFD points. As this ratio decreases, the optimizer has more and more control over the individual CFD points and is likely to produce a non-smooth surface. A ratio of at least four CFD points to one FFD point is recommended. The distribution of the control points in relation to the wing is shown in Figure 6.

Both the aerodynamic surface mesh and the structural mesh points are embedded in the FFD volume. The control points of the FFD volume are used to create a trivariate B-spline mapping of its interior. This mapping defines the parametric position of each node of the embedded surface mesh and structural mesh. As the positions of the FFD control points change, the embedded geometry deforms continuously according to the B-spline mapping. The design variables are set up to manipulate the position of the control points to enact local or global changes to the shape of the embedded geometry. Deformations of the surface mesh are then propagated out to the volume mesh by the mesh-warping algorithm. The derivatives of the volume mesh nodes with respect to the surface mesh nodes are

computed using automatic differentiation. The derivatives of the surface mesh nodes with respect to the FFD control points are computed analytically, and the derivatives of the FFD control points with respect to the design variables are computed using the complex-step method [9].

Our parametrization uses both global and local design variables. The global design variables act on a group of the FFD control points, facilitating large-scale deformations. Thus, a given global design variable produces a nonzero derivative for multiple control points. The local shape design variables control the individual displacement of each control point. It follows that the Jacobian of the control points with respect to the local design variables is the identity matrix. When multiple design variables affect a given control point, the operations are combined linearly. Figure 6 demonstrates the use of global taper and sweep design variables to convert the rectangular wing to a planform resembling that of the CRM wing.

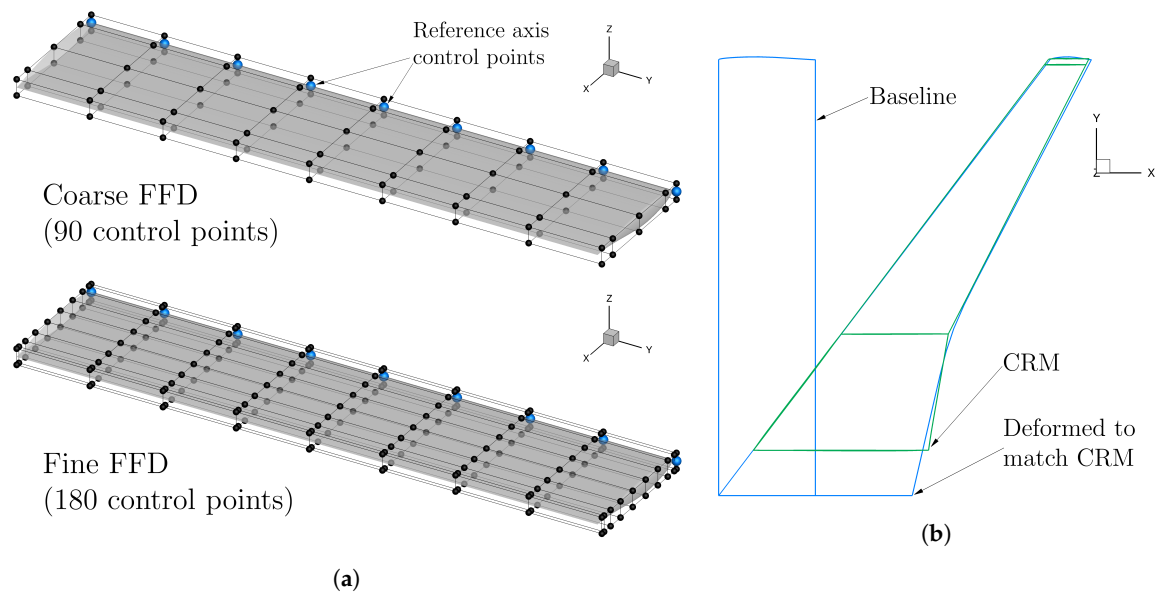


Figure 6. (a) Wing surface embedded in the both coarse and fine FFDs. (b) The CRM planform is reproduced by modifying the FFD control points.

2.3. Optimization Problem

The optimization problem is defined in Table 4. Bounds and scaling factors for the variables and constraints are listed in the columns to the left, where applicable. The optimization problem requires three high-fidelity analyses: (1) nominal cruise, (2) 2.5 g pull-up maneuver, and (3) 1.3 g cruise buffet. The flow conditions for these three cases are listed in Table 5. The thermodynamic state of the freestream for each flight condition is determined from the Mach number and altitude specified in Table 5 in conjunction with the International Standard Atmosphere model [50].

Table 4. Rectangular wing aerostructural optimization problem description.

		Quantity	Lower	Upper	Scaling
minimize	W_{fuel}	1			
with respect to	Angle of attack	3	0°	10°	0.1
	Twist	7	−10°	10°	0.05
	Sweep	1	0 m	25 m	0.01
	Chord scaling	3	0.25	2.0	0.1
	Sectional shape	180	−50 cm	50 cm	1
	Panel thickness	131	2 mm	20 cm	100
	Stiffener thickness	108	2 mm	20 cm	100
	Stiffener height	91	5 mm	10 cm	100
	Stiffener pitch	3	10 cm	30 cm	100
	Panel length	108			
	Fuel tractions	301			
Total fuel mass	1				
Total number of design variables		937			
subject to					
<i>Nonlinear constraints</i>	$L_{\text{nominal}} = (W_2 + W_3)/2$ $L_{2.5g} = 2.5W_2$ $C_{L,\text{buffet}} = 1.3(C_{L,\text{nominal}} + 0.05)$ Structural failure constraints Buffet-onset constraint $S_{\text{ref}} - S_{\text{ref,orig}}$ Minimum wingtip thickness Minimum trailing edge thickness Minimum spar height thickness Total fuel mass constraint Fuel volume constraint Fuel traction consistency constraints Panel length consistency constraints LE/TE constraints Monotonic constraint on chord scaling	1	0	0	10^{-6}
		1	0	0	10^{-6}
		1	0	0	10^{-6}
		5		1	1
		1		0.04	100
		1	0	0	0.1
		15	10%		1
		15	100%		1
		30	60%		1
		1			
		1			
<i>Linear constraints</i>	$t_{\text{stiff},i} - t_{\text{panel},i}$ $h_{\text{stiff},i} - t_{\text{stiff},i}$ $t_{\text{panel},i} - t_{\text{panel},i+1}$ $t_{\text{stiff},i} - t_{\text{stiff},i+1}$ $h_{\text{stiff},i} - h_{\text{stiff},i+1}$	301			
		108			
		18			
		2			
		108	−2 mm	2 mm	
		108	0		
		104	−2.5 mm	2.5 mm	
104	−2.5 mm	2.5 mm			
88	−5 mm	5 mm			
Total number of design constraints		1013			

Table 5. Flow conditions for high-fidelity analyses.

Case	Mach	Altitude (ft)	Re
Nominal cruise	0.85	37,000	37.7×10^6
2.5 g pull-up maneuver	0.64	0	91.2×10^6
1.3 g cruise buffet	0.85	37,000	37.7×10^6

2.3.1. Objective Function

The objective function is calculated using Equation (1). As mentioned previously, the optimizer can change W_{wing} , C_L , and C_D to decrease W_{fuel} . The wing mass is made up of the mass of the finite-element model and an additional component to account for any increase in wing area, which is given by

$$W_{\text{wing}} = 2.5 W_{\text{wingbox}} + 4000 \frac{S_{\text{ref}}}{S_{\text{ref,orig}}}. \quad (2)$$

To account for the lack of the fuselage and other surfaces, we make the following modifications to the lift and drag coefficients calculated by the CFD:

$$C_{L,\text{total}} = \frac{2.2 L_{\text{wing}}}{S_{\text{ref}} q}, \quad (3)$$

$$C_{D,\text{total}} = \frac{2 D_{\text{wing}} + 3.453q}{S_{\text{ref}} q}, \quad (4)$$

where L_{wing} and D_{wing} correspond to a RANS solution of a half-wing CFD model, thus the factor of two in Equations (3) and (4). The drag area of 3.453 m^2 in Equation (4) is the product of the baseline reference area and the sum of the drag coefficients for the fuselage, empennage, and nacelle surfaces. The nacelle and vertical stabilizer combined contribute 30 drag counts; the same value used in the uCRM optimizations. To calculate the drag markup for the fuselage and horizontal stabilizer, we first compute the ratio of the mesh-converged drag values from the wing-only CRM case [21] and the DPW 4 wing-body-tail geometry. This ratio is then used to calculate the fraction of the uCRM drag corresponding to the fuselage and horizontal stabilizer. The complete calculation is as follows:

$$C_{D,\text{FH}} = \left(1 - \frac{C_{D,W,\text{CRM}}}{C_{D,\text{WFH},\text{DPW4}}}\right) C_{D,\text{WFH},\text{uCRM}} \approx 0.006, \quad (5)$$

where W, F, and H refer to wing, fuselage, and horizontal stabilizer, respectively. These modifications are used for all the high-fidelity analyses.

2.3.2. Design Variables

Each case has an angle-of-attack variable to enable matching the lift constraint. There are seven twist variables, each controlling the rotation of one of the spanwise FFD sections about the leading edge. The first two sections are fixed at zero twist. The sweep variable corresponds to the streamwise displacement of the wingtip leading edge. All other FFD sections are displaced linearly to create a straight leading edge. Chord scaling is controlled at FFD sections 1 (symmetry), 4, and 9. The intervening spanwise sections are scaled linearly to ensure straight leading and trailing edges. An example of the sweep and chord scaling variables is shown in Figure 6. As explained in Section 2.2.3, we use an FFD with 90 control points for the coarse CFD meshes and an FFD with 180 control points for the fine meshes. At each spanwise section, the FFD control points are restricted to in-plane displacements that are perpendicular to the freestream. The FFD control points regulate the cross-sectional shape of the wing.

The structure is divided into 131 design variable groups: 23 for the ribs, 18 for the spars, and 45 groups each for the upper and lower skins. Each of the rib design variable groups has a single panel thickness variable because the ribs are not modeled with the smeared stiffeners. The spar design variable groups share a single variable for stiffener height and another for stiffener pitch. Each spar group has its own variables for panel thickness and stiffener thickness. A single stiffener pitch variable is shared by all upper skin groups, and another is shared by all lower skin groups. All skin design variable groups have their own variables for panel thickness, stiffener thickness, and stiffener height. Finally, each of the design variable groups with smeared stiffeners (skins and spars) has a variable for panel length. The alignment of the stiffeners on the panels is calculated based on the initial panel reference axis, but does not change in the course of the optimization. Thus, while the stiffness matrix is updated to reflect changes in the length of the stiffeners, it does not account for changes in the orientation of the stiffeners. For optimization problems that allow the wing sweep to change, the stiffeners remain aligned with the initial sweep. This creates an artificial benefit for the wing sweep to remain close to the initial value, because the structure is most efficient when the stiffeners are aligned with the wing sweep. For small variations in wing sweep, the effect of this discrepancy is minor; however, for the rectangular wing case, we expect large changes in wing sweep. We have two ways of managing this issue. First, for optimization problems that start from the baseline rectangular wing, we manually set the stiffener orientation to a value that is close to the expected optimal wing sweep. Second, we use a multi-level approach to optimization, in which successive optimizations start from where the previous one left off (see Section 3.1).

The weight of the fuel in the wing is applied as a uniform traction to the lower skin of the wingbox between the symmetry plane and the 44th rib. This region of the lower skin is made up of 301 panels,

each of which is associated with a design variable for the weight of fuel to be applied as a traction. Additionally, there is a variable for the total amount of fuel in each case. The cruise and buffet cases have a fuel load corresponding to the mid-cruise point, while the maneuver case has a fuel load corresponding to the full fuel load.

2.3.3. Design Constraints

Each of the analysis points has its own lift constraint. The structural failure constraints are computed by aggregating a failure criterion over a group of elements using the Kreisselmeier–Steinhauser (KS) function [51]. There are two types of failure criterion added as constraints. The stress failure criterion is calculated as the von Mises stress in an element multiplied by a safety factor and divided by the material yield strength. The buckling failure criterion is based on the critical buckling load of a stiffened quadrilateral panel. We include three stress failure constraints: one for the ribs and spars, one for the upper skin, and one for the lower skin. Two buckling failure constraints are added: one for the ribs and spars, and one for the upper skin. The buffet-onset constraint was developed and validated by Kenway and Martins [52]. This constraint is computed by a correlation of buffet to the amount of separated flow on the upper surface of the wing. This constraint is added to the 1.3 g cruise case, to maintain the required 30% margin to buffet during cruise.

A constraint is added to preserve the original reference area. Thickness constraints are added along the wingtip to prevent excessive flattening of the wingtip cap, which could crush cells in the volume mesh. These constraints limit the thickness of the wingtip cap to a minimum of 10% the original thickness. Any decrease of the trailing edge thickness is prevented with a set of 15 thickness constraints at 99% chord. Thickness constraints are also added along the fore and aft spars to prevent decrease beyond 60% of the original value. This set of thickness constraints was added to prevent excessive thinning of the outboard wing (see Figure 11). A set of constraints (dubbed LE/TE constraints) are added to the pairs of FFD control points at the leading and trailing edges of each section to ensure equal and opposite displacement. This ensures that the shape variables do not cause shearing twist, which would be redundant with the global twist variables. Bons et al. [26] used a monotonic constraint on chord variables to ensure that the chord decreases monotonically from the root to the wingtip. In this work, we include a linear constraint on the chord scaling variables to enforce this property.

There are two fuel load constraints to ensure that the total fuel load variable for each case is consistent with the actual amount of fuel being carried by the aircraft. Each of the design variable groups with smeared stiffeners has a linear constraint to maintain a difference of less than 2 mm between the panel thickness variable and the stiffener thickness variable. There are also linear adjacency constraints to limit the difference in stiffener height, stiffener thickness, and panel thickness between adjacent panels to 1 cm, 5 mm, and 5 mm, respectively. There are 64 nonlinear constraints added to ensure the deformed panel lengths are consistent with the 64 panel length variables. An additional 602 nonlinear constraints exist to ensure that the fuel traction variables are set to the correct value. Finally, for each fuel load, there is a volume constraint to ensure that the fuel can fit inside the wingbox.

2.4. Structural Pre-Optimization

Initially, all the structural members have uniform thickness and stiffener sizings. We can start the aerostructural optimization from a more reasonably sized structure if we first optimize the structural sizing with a set of fixed aerodynamic loads. This structural pre-optimization is a cycle with five iterations. In each iteration, we run an aerostructural analysis with the current structural sizing to get the aerodynamic loads. Then, we apply the aerodynamic loads to the structure and run an optimization that minimizes the structural mass with respect to the failure constraints on the 2.5 g maneuver. We repeat this process five times to arrive at a semi-converged aerostructural state. For the plank geometry, the structural pre-optimization produces a wingbox weight (W_{wingbox}) of 11,874 kg, which corresponds to a total wing weight (W_{wing}) of 37,686 kg. The optimized structure for the CRM-shaped planform has a wingbox weight of 10,968 kg and a total wing weight of 35,954 kg.

2.5. Multipoint Optimization Problem

Sophisticated methods for determining a representative set of design conditions and weights have been researched by Toal and Keane [53] and by Liem et al. [37]. However, in this work we use a simple 3-point stencil in C_L -space based on the ADODG Case 4.2 [30]. The nominal C_L is based on the mid-cruise weight at Mach 0.85 and 37,000 ft, just as in the single-point optimization problem. The auxiliary design points are analyzed at $C_{L,nominal} \pm 0.05$ at the same Mach number and cruise altitude. Initially, we tried $C_{L,nominal} \pm 0.025$, following Brooks et al. [18], but we found little variation between the single-point and multipoint designs. The three design points are weighted equally, so the objective function is the arithmetic mean of the fuel burn calculated from L/D at each flow condition. In all other respects, the multipoint optimization is identical to the problem described in Table 4.

2.6. Low-Speed, High-Lift Separation Constraint

We also experiment with a novel approach for ensuring airworthiness at low-speed, high-lift conditions. Using publicly available flight data of the Boeing 777-200ER (an aircraft with similar specifications to the CRM), we determined Mach 0.4 and 10,000 ft to be a low-speed flight condition that should exhibit good aerodynamic performance. This flight condition is shown in relation to the flight data in Figure 7. As will be shown, both the single-point and multipoint designs exhibit severe separated flow when evaluated at this flight condition at a high angle of attack. By contrast, a single-point optimization with fixed RAE 2822 cross-sections performs well.

Instead of minimizing the drag at this flight condition, we use a constraint to limit separation on the wing. Our approach is inspired by buffet-onset constraint of Kenway and Martins [52] and uses the same formulation for the separation sensor. Based on the analysis of the single and multipoint wings, the constraint allows no more than 10% of the upper surface of the wing to have separated flow. The analysis point for the separation evaluation is constrained to generate enough lift to sustain the nominal takeoff weight (W_2) of the aircraft at Mach 0.4 and 10,000 ft. Both the chosen flight condition and the value for the upper limit of the separation constraint are somewhat arbitrary and are subject to change for other applications. This work is mainly concerned with introducing the application of a low-speed separation constraint and evaluating its impact on the optimized design. The correlation of this constraint with established airworthiness regulations remains for future work.

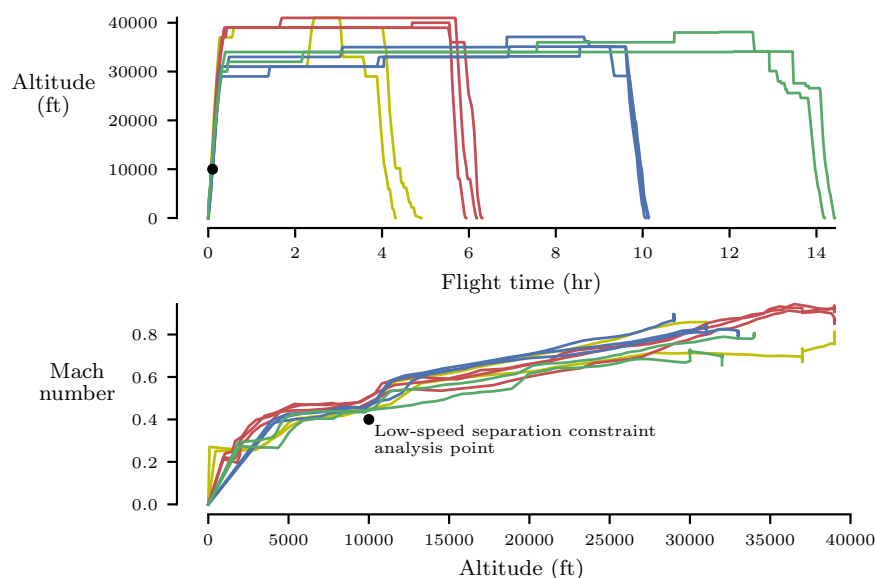


Figure 7. Analysis point for low-speed, high-lift separation constraint is placed at the boundary of the climb profile for the Boeing 777-200ER.

3. Results

3.1. Multi-Level Optimization Procedure

In previous work on aerodynamic shape optimization by Lyu et al. [21], a sequence of optimizations was performed on progressively finer meshes to reduce the total computational time of the optimization. Each optimization starts where the previous optimization ends, so that the maximum benefit is extracted from each mesh. The coarser meshes do not resolve the flow as accurately, but they still provide derivatives that point the optimizer in the direction of the true optimum. With each successive mesh level, the design is closer to the optimum with an enriched set of derivatives to guide its path. For this study, we adopt the multi-level optimization method with the three coarsest meshes: B2, A2, and B1. The purpose of this study is to compare the multi-level approach to an optimization using only the B1 mesh. The optimization problem is a simplified version of the one laid out in Table 4, where the sweep is fixed at 34.9° . Thickness constraints along the spars are not included in this comparison.

The result of this comparison is shown in Figure 8. Both methods converge to essentially the same shape, but the computational cost of the multi-level approach is 60% of the cost of the single-level optimization. The direct comparison does not tell the whole story. In practice, the process of setting up and troubleshooting a new optimization problem requires countless debugging runs. We have found great value in having a very coarse mesh during this initial phase. The cost savings due to troubleshooting with a coarse mesh are not easily accounted for, but far outweigh the cost savings due to a single optimization run.

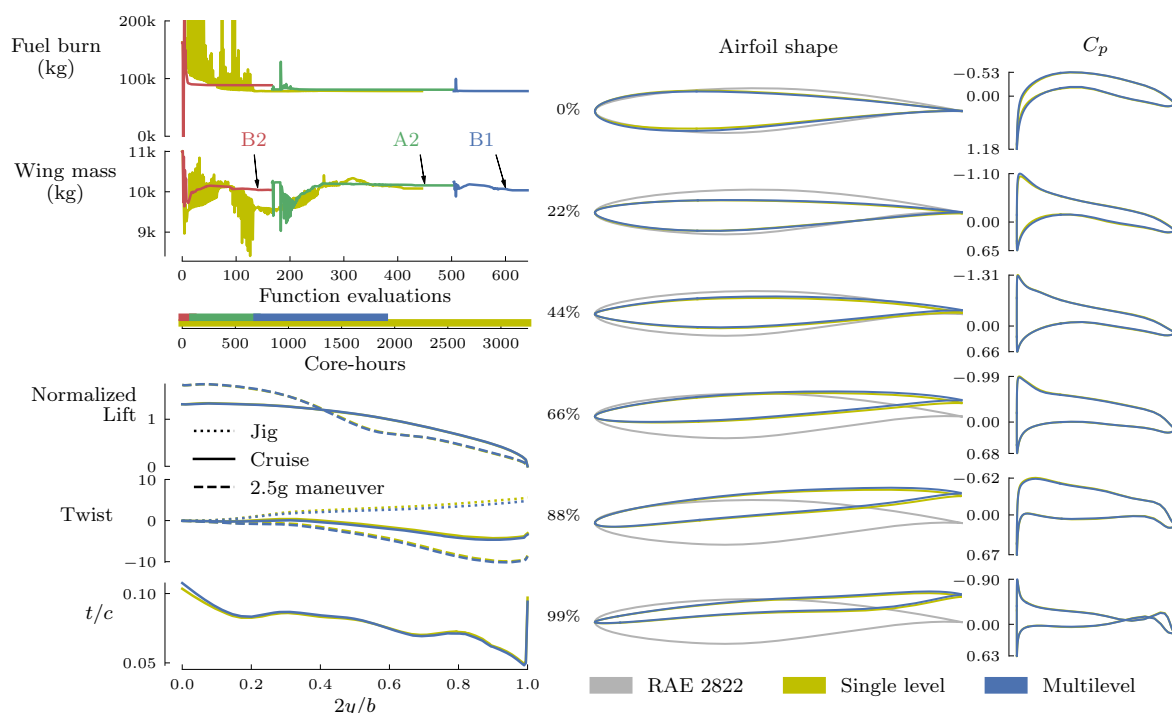


Figure 8. For this problem, the multi-level approach achieves essentially the same design as the single-level optimization with the finest mesh—at 60% of the computational cost.

3.2. Single-Point Optimization

Now we look at the full optimization problem described in Table 4. For this problem, we use mesh levels B2, A2, B1, and A1 successively to arrive at the final result. The coarse FFD (90 control points) is used on mesh levels B2 and A2 and the fine FFD (180 control points) is used on mesh levels

B1 and A1. The optimized wings from each mesh level are shown in Figure 9. The planform of the starting geometry is shown in gray.

The transonic flow condition presents an interesting trade-off between wave drag and wing mass in relation to the sweep variable. To decrease wave drag, the optimizer can increase wing sweep or modify airfoil shape. However, increasing the wing sweep results in a heavier wing to support increased bending loads. For the B2 wing, the tip sweep reaches the upper limit. The airfoils of the A2 wing are very similar to the B2 wing, and yet the sweep decreases by 2° , indicating that the improved resolution of the flow field favors less sweep. The final two optimizations converge to a sweep value in between the first two. Switching to the fine FFD allows the optimizer to fine-tune the airfoil shape for the single design point, resulting in a fairly sharp leading edge and a pronounced suction peak. A distinctive lower surface concavity forms at the leading edge, which is reminiscent of the result of an airfoil optimization of the RAE 2822 by Drela [29]. The optimizations that use the finer meshes are also able to produce more passive load alleviation, as shown in the difference between the load and twist distribution at the cruise and maneuver design points.

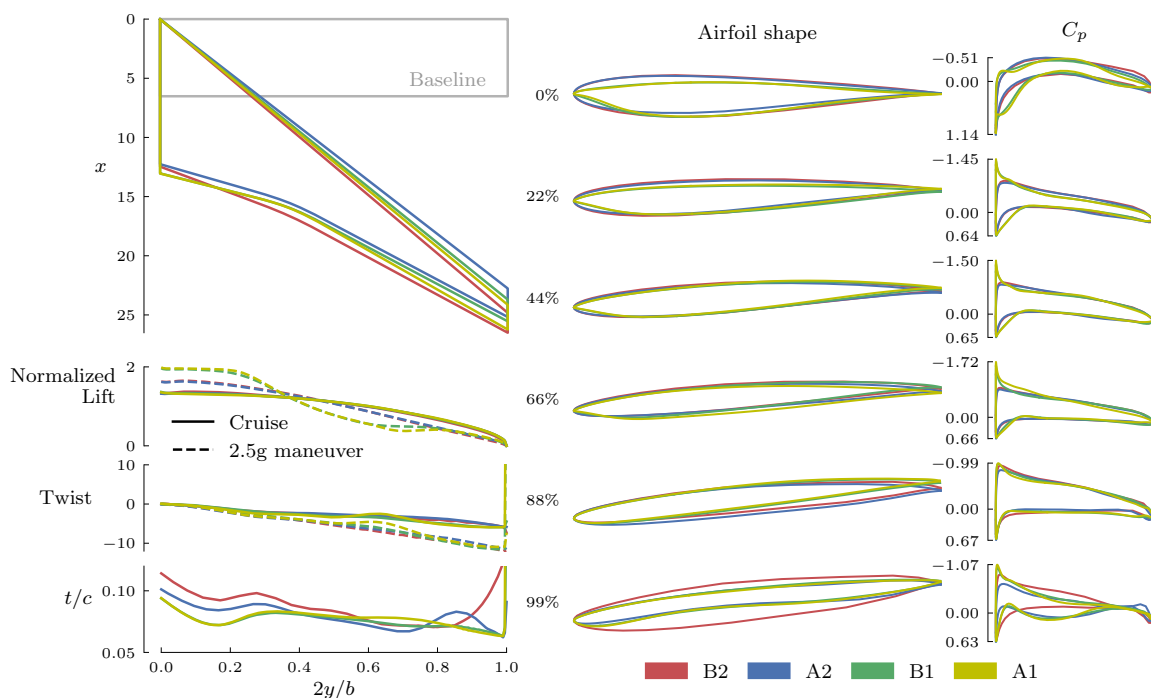


Figure 9. Each successive level of mesh refinement yields additional design changes.

Figure 10 shows that the improvements achieved from each successive optimization are preserved in a mesh refinement study. The optimized wings from each of the four mesh levels are analyzed with the finer meshes. The flow condition for this comparison is the cruise condition at $C_L = 0.5$. We chose $C_L = 0.5$ because each of the wings was optimized to a lift coefficient near that value. The baseline mesh convergence is also plotted to show the increase in drag as the large shock structure is captured more accurately on successively finer meshes. This comparison gives an indication of the value added for each successive optimization in the multi-level approach. Compared to the B2 optimum, meaningful gains are realized in the A2 and B1 optimizations. However, the drag reduction from the B1 optimum to the A1 optimum is marginal. The minimal differences in the wing design between the B1 and A1 optima suggest that there is no need to continue the multi-level optimization onto the next mesh level.




Figure 10. Drag convergence study for the baseline and single-point optimized wings. For the baseline wing, the drag increases as the mesh is refined because the shock is resolved more accurately on finer meshes. The improvements to the optimized wings are preserved as the mesh is refined.

In addition to starting the optimization with the plank geometry, we also ran the optimization using the CRM planform as the initial design. In addition to the differences in starting planform, the initial structural sizing for each of these runs differs because of the structural pre-optimization (see Section 2.4). As shown in Figure 11, the optimization converges to nearly the same design starting from the plank geometry and starting from the CRM planform. A closer look at the differences in final design variables reveals that while the optimization did converge to the same general shape, there are significant differences between the two optimized designs, especially in the structural sizing (Figure 12). Notably, the design variables of larger magnitude are more likely to converge to the same value, whereas smaller design variables exhibit greater variance.

Figure 11 also shows the effect of the number of shape variables on the overall design. When the coarse FFD is used, the wing has reduced sweep and is unable to produce as much passive load alleviation as the wing optimized with the fine FFD. Additionally, the optimizer is unable to tailor the leading-edge radius as precisely with the coarser FFD. The wing optimized without shape variables converges to a planform design with a constant chord on the inboard section of the wing. Normally, the optimizer can reduce wave drag by increasing chord, while keeping the thickness constant, but in this case, the optimizer avoids increasing the root chord because it has no control over the thickness ratio of the wing. As shown in Table 6, including shape variables in the optimization reduces the objective by nearly 14,000 kg. Further fuel burn reduction of over 1000 kg is realized by using the fine FFD over the coarse FFD. This improvement is likely due more to the 2000 kg reduction in wing weight than the marginal improvement in L/D .

Table 6. Results of single-point optimization.

Case	Mesh Level	FFD	W_{fuel} (kg)	W_{wing} (kg)	L/D	Sweep (deg)
Starting from plank	B2	Coarse	88,740	37,531	20.3	37.3
	A2	Coarse	80,727	35,758	21.7	35.2
	B1	Fine	77,243	32,725	22.3	36.0
	A1	Fine	75,834	31,919	22.5	36.7
Starting from CRM	A1	Fine	75,667	31,002	22.5	36.7
	Coarse FFD	Coarse	76,941	33,027	22.4	33.1
No shape variables	A1	Fine	90,928	39,167	20.0	38.1

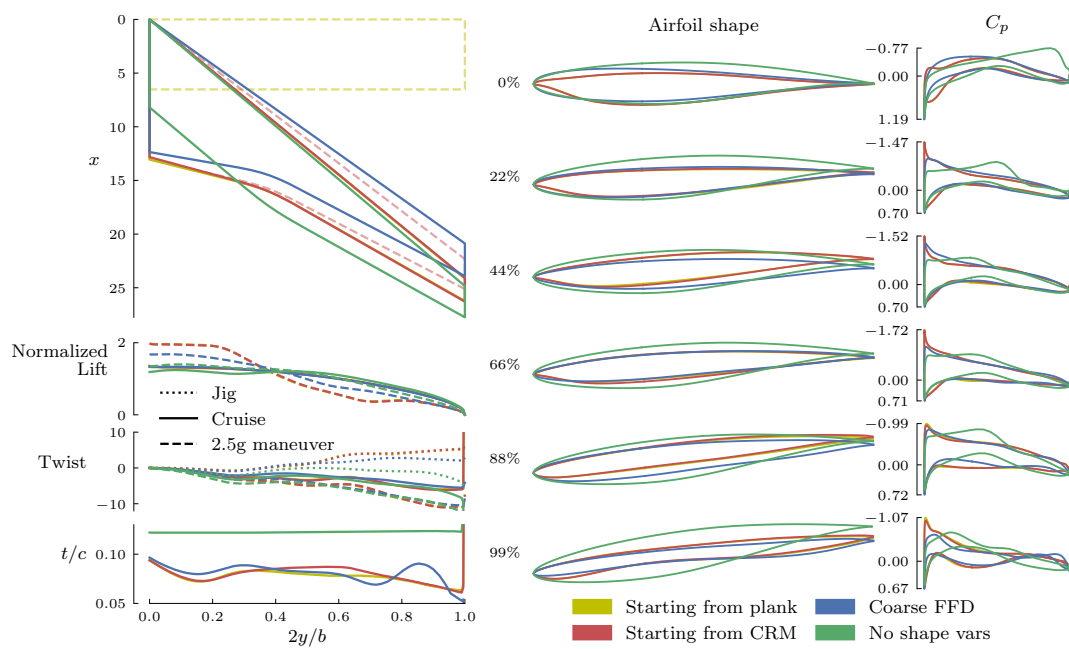


Figure 11. Comparison of single-point optimized designs. The optimizations starting from the plank and CRM planforms converge to nearly the same design.

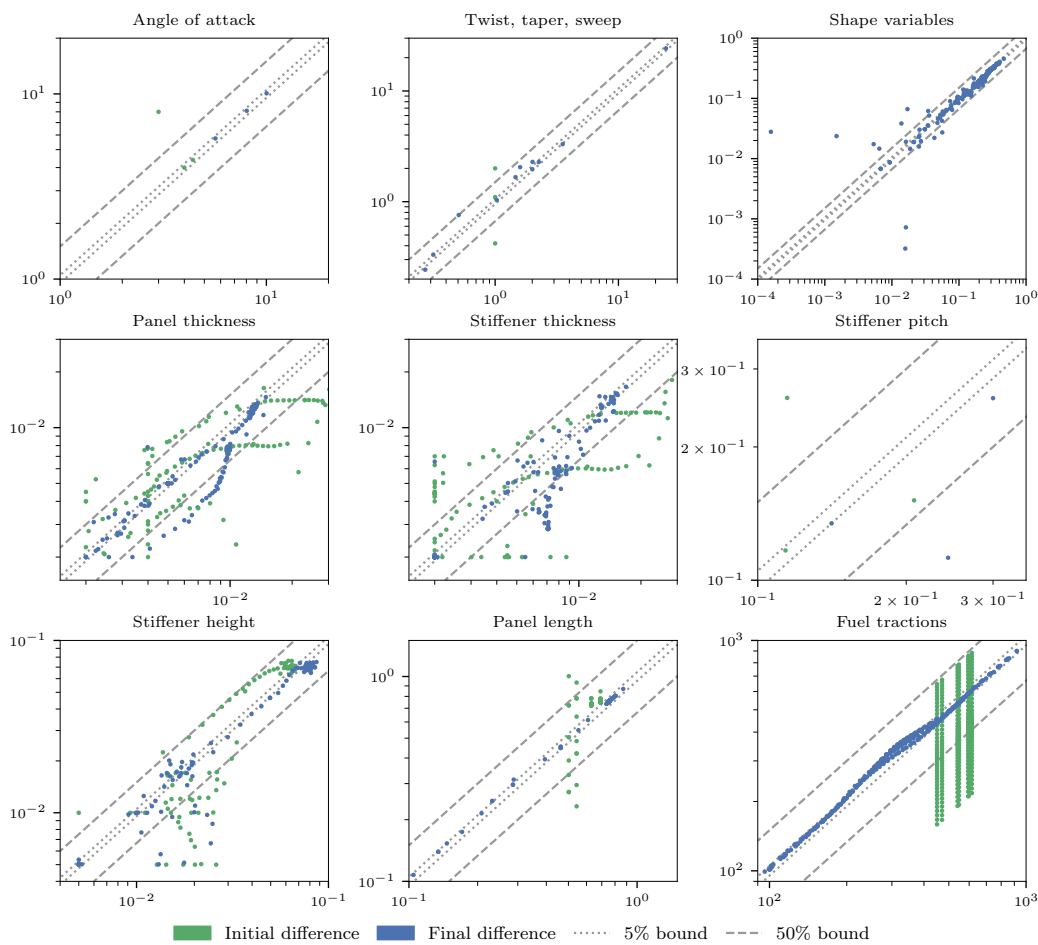


Figure 12. The difference in initial and final design variables for wing optimization problems starting from a plank planform and a CRM planform.

3.3. Robust Design Optimization

The main objective of this study is to evaluate the impact of two different robust optimization approaches: (1) a multipoint objective and (2) a separation constraint at a low-speed, high-lift flight condition. The designs considered in this section are:

- (SP) The single-point design from Section 3.2 (started from the CRM planform).
- (MP) A three-point design with $C_{L,nominal}$, $C_{L,nominal} - 0.05$, and $C_{L,nominal} + 0.05$ at Mach 0.85 and 37,000 ft.
- (SP-LS) A single-point design with a separation constraint at Mach 0.4 and 10,000 ft.
- (SP-NSV) Result of single-point optimization without shape variables (RAE 2822 cross-sections).

Although only the MP optimization problem included the multipoint stencil in the objective function, the other three optimized designs were analyzed at the multipoint flight conditions for comparison. Characteristics of the four optimized designs are listed in Table 7. The multipoint design achieves the best average fuel burn, but does slightly worse than the single-point design at the nominal flight condition. The SP-LS design burns more fuel across the three points than the SP (+1.0%) and MP (+1.6%) designs, but compared to the SP-NSV design it is 16.0% more fuel efficient. The relative performance of the four designs is shown in Figure 13.

Surprisingly, the airfoils of the multipoint design are nearly identical to the single-point design (Figure 14). Increased sweep on the multipoint design is the major geometric difference between the two. Accordingly, the multipoint wing is slightly heavier than the single-point wing, with most of the weight gain in the skins and the aft spar. By contrast, the SP-LS wing is substantially different than the single-point or multipoint wings.

Table 7. Optimization results.

Property	Units	SP-NSV	SP	MP	SP-LS
Combined fuel burn	kg	92,145	76,592	76,261	77,518
W_3	kg	188,167	180,001	180,191	183,002
Wing weight	kg	39,167	31,001	31,191	34,001
Upper skin	kg	5040	3957	3976	4529
Lower skin	kg	4902	3643	3637	4238
Ribs	kg	1691	1234	1260	1235
Fore spar	kg	409	165	187	217
Aft spar	kg	423	199	214	180
Sweep	deg	38.1	36.7	37.4	30.8
<i>Nominal</i>					
Fuel burn	kg	90,928	75,667	75,817	77,518
Angle of attack	deg	5.57	5.77	5.65	4.70
C_L		0.5443	0.5075	0.5082	0.5157
C_D	counts	272.3	226.0	226.4	229.1
L/D		19.99	22.46	22.44	22.51
$C_{L,nominal} - 0.05$					
Fuel burn	kg	91,642	79,147	78,972	80,074
Angle of attack	deg	4.92	5.19	5.06	4.24
C_L		0.4938	0.4575	0.4582	0.4654
C_D	counts	248.6	211.6	211.3	214.3
L/D		19.86	21.63	21.69	21.72
$C_{L,nominal} + 0.05$					
Fuel burn	kg	93,866	74,963	73,995	75,765
Angle of attack	deg	6.26	6.33	6.22	5.15
C_L		0.5937	0.5575	0.5581	0.5657
C_D	counts	304.9	246.3	243.6	248.6
L/D		19.48	22.64	22.91	22.75




Figure 13. Optimizing for robust performance in cruise and climb incurs ~1000 kg increase in cruise fuel burn.




Figure 14. The multipoint design is very similar to the single-point design, but including the low-speed separation constraint elicits striking modifications.

One of the most striking differences is the disappearance of the distinctive concavity at the leading edge of the lower surface. The SP-LS cross-sections have a larger leading-edge radius as a result and seem more typical of a traditional airfoil. The C_p curves feature reduced suction peaks (owing to larger leading-edge curvature) and more aft loading than the other designs. Sweep is reduced by 5° compared with the single-point design. The distribution of t/c is significantly lower from 40–80% span, likely decreasing wave drag, which would otherwise increase because of the reduction in sweep. Passive load alleviation at the 2.5 g maneuver condition is severely degraded for the SP-LS wing.

The variation in spanwise loading between the cruise and maneuver conditions is minimal, as opposed to the SP and MP designs where the cruise loading is elliptical and the maneuver loading is bell-shaped. This means that the SP-LS wing requires a heavier structure to achieve an elliptical cruise lift distribution and still satisfy the failure constraints at the maneuver condition. The upper and lower skins see the greatest increase in weight, but the fore spar is also significantly heavier than the single-point design.

Given the differences in weight and geometry, the cruise drag polars for the SP, MP, and SP-LS designs are surprisingly similar (Figure 15). These three designs have nearly the same performance at the nominal cruise point; the variation in nominal fuel burn is due to the differences in weight rather than the aerodynamic efficiency. Moving outward from the nominal design point in either

direction, the multipoint design is the most robust, followed by the SP-LS design. The multipoint design does especially well as C_L is increased from the nominal design point. Most of the improvement in the average fuel burn over the single-point design comes from the $C_{L,nominal} + 0.05$ design point. The SP-LS design achieves higher L/D than the SP design at all three design points yet burns more fuel because of the heavier structure.

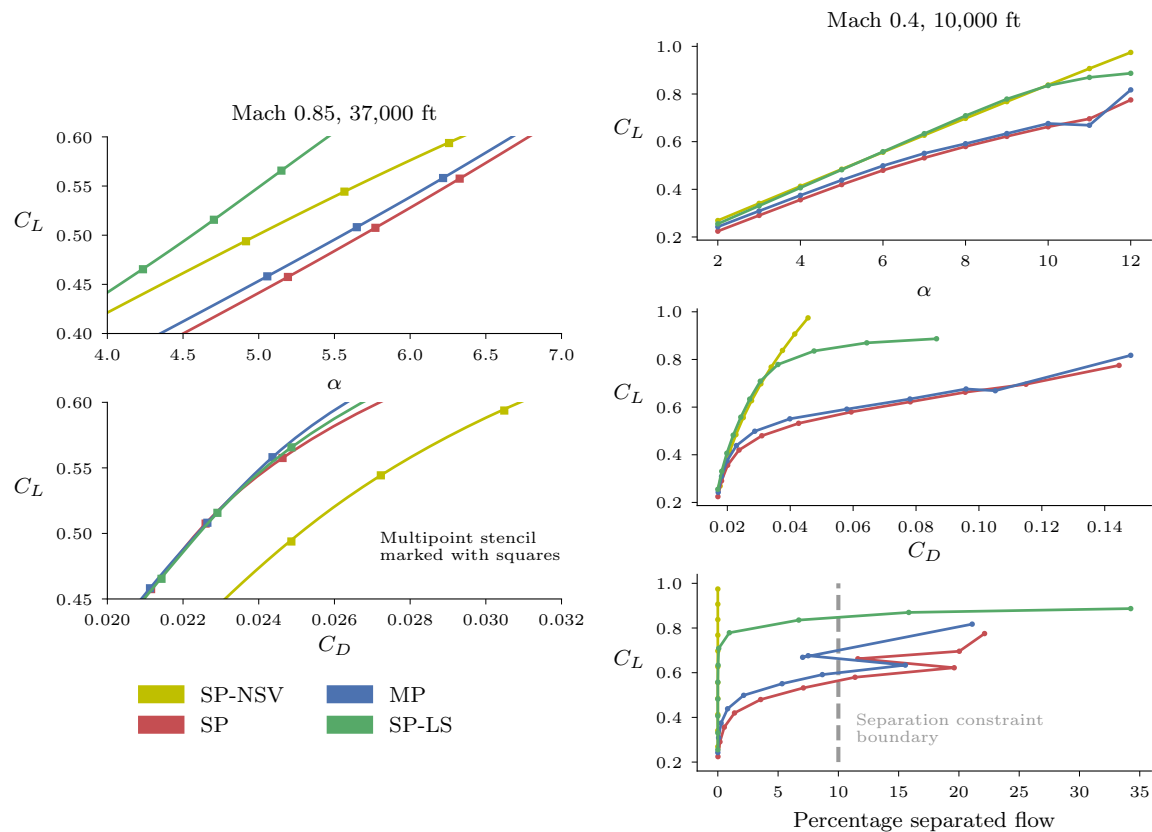


Figure 15. The low-speed separation constraint improves robustness at both cruise and climb flight conditions.

The right side of Figure 15 shows the performance of the four designs at the climb flight condition. The four designs were analyzed across a sweep of angle of attack at increments of 1° . The SP and MP designs do not converge when the angle of attack exceeds 9° due to massive separation on the upper surface. Separated flow leads to a sharp increase in C_D as angle of attack is increased for both designs. The trend of the drag polars indicates that the $C_{L,max}$ for these wings would be lower than the C_L required to satisfy the lift constraint $L = W_2$.

The multipoint design fares slightly better than the single-point design, but both fail to meet expected airworthiness at the climb condition. The drag polar of the SP-LS wing mimics that of the SP-NSV wing up until 9° angle of attack. At that point separation ensues, but the 10% threshold is not exceeded until nearly 11° angle of attack. At 9° angle of attack, the SP-LS wing generates 24% more lift and 54% less drag than the SP wing. The near-complete elimination of on the SP-LS wing at 9° angle of attack is notable (Figure 16).

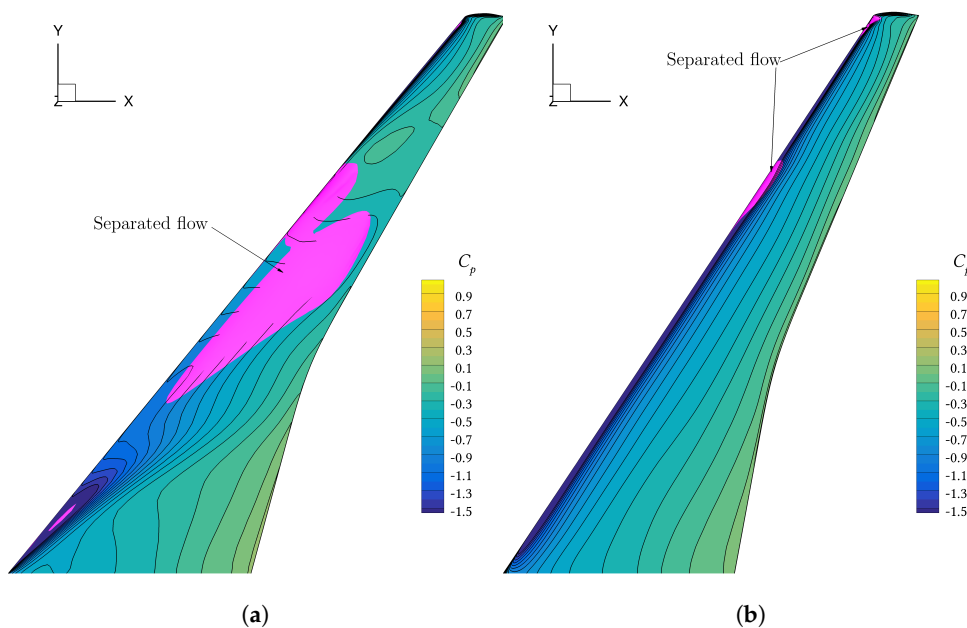


Figure 16. Separation on the upper surface at 9 degrees angle of attack is nearly eliminated with the low-speed separation constraint. (a) Single-point. (b) Single-point with low-speed separation constraint.

The buffet-onset constraint is inactive at the end of the SP-LS optimization, which suggests that low-speed separation is a more stringent requirement than buffet-onset. Conversely, the buffet-onset constraint was active for both the single-point and multipoint designs, but did not improve the low-speed, high-lift performance of these wings. One possible reason for this finding is that the low-speed separation constraint has the effect of decreasing the cruise angle of attack. Separation at the 1.3 g flight condition is more likely to occur if the wing is cruising at a higher angle of attack, as is the case for the SP and MP designs.

4. Discussion

In Section 3.2, we investigated several aspects of the single-point aerostructural optimization problem. We compared results obtained using various CFD mesh levels and saw that the final cross-sectional shape, planform, and structural sizing are all affected by the mesh level. This comparison also revealed that there are diminishing returns when using increasingly finer CFD meshes. We showed that the optimizer was able to converge to a similar design from two radically different starting points, allaying concerns that gradient-based optimization is not suitable for design space exploration.

However, we found that the structural sizing variables did not converge as closely to the same values as the variables for the wing shape. This finding is consistent with our experience that aerostructural optimizations are more difficult to converge and do not converge as tightly as aerodynamic-only optimizations. This is in part because of the high condition number of the stiffness matrix for shell finite elements, which limits the achievable numerical precision for both structural and aerostructural analyses and gradient computations. Additionally, the KS functions used to aggregate structural failure criteria are highly nonlinear and can create problems for gradient-based optimizers. More work is needed to diagnose and eliminate these issues with aerostructural optimization convergence. Finally, we considered the impact of the number of shape variables on the overall design and found that the airfoil shape is implicitly linked to the planform shape and the structural sizing.

Multipoint objective functions typically consider design points in the cruise regime. Although this does result in more robust cruise performance, we have shown that cruise-optimized designs do not, in general, perform well at other conditions, such as climb. We have shown that the inclusion of a

low-speed separation constraint in the SP-LS case pushes the optimization to a completely different design than the single-point and multipoint problems. The single-point, multipoint, and SP-LS designs do not differ significantly in terms of aerodynamic performance at cruise. Rather, the changes in the wing design seem to have the greatest impact on performance at off-design conditions.

The low-speed separation constraint dramatically improves the high-lift capability of the wing at low-speed conditions. Additionally, the design changes required to satisfy the separation constraint have the effect of reducing passive load alleviation at the 2.5 g maneuver loading, resulting in a substantial increase in structural weight.

The differences in fuel burn among the designs are more related to varying structural weights than significant stratification of cruise performance. The single-point and multiple designs are not viable concepts because of their poor low-speed, high-lift characteristics, whereas the SP-LS is a much more practical design.

5. Conclusions

The primary purpose of this work is to demonstrate that MDO is an effective means of exploring the aerostructural wing design space. We showed that an optimizer could start from a rectangular, constant cross-section wing and traverse the design space to arrive at an optimized transonic swept wing with custom airfoils and optimally sized structure. Moreover, the optimizer arrived at the same design (within a small tolerance) when the optimization began from a swept planform. The results show no evidence of multiple local minima in the OML shape for this design problem. Although some of the structural design variables showed significant differences, this is most likely due to difficulties in converging the aerostructural optimization beyond a certain tolerance rather than to some physically significant multimodality. In any case, the differences in the structural design variables did not prevent the OML design variables from converging to the same values. To reduce the computation cost of these studies, we used a multi-level optimization process, in which the CFD meshes are refined in successive optimization runs, to reduce the overall computational cost by 40%.

This work also introduces a novel method to improve off-design robustness in optimized wings. First, we showed that both single-point and multipoint cruise-optimized designs exhibit massive separation at a low-speed, high-lift flight condition representative of a typical climb profile. A separation constraint applied at the climb condition restores attached flow without severely degrading cruise performance. In wing design optimization, leading edge thickness constraints are usually used to prevent the optimization from excessively reducing the leading-edge radius.

With the proposed low-speed separation constraint, adequate curvature was preserved on the leading edge without having to resort to thickness constraints. This underscores the importance of considering off-design performance into the optimization problem when designing wings.

Author Contributions: Conceptualization, methodology, writing—review and editing, visualization, N.P.B. and J.R.R.A.M.; software, validation, formal analysis, investigation, data curation, writing—original draft preparation, N.P.B.; supervision, project administration, funding acquisition, J.R.R.A.M. All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded by the University of Michigan Rackham Merit Fellowship Program.

Acknowledgments: The authors would like to thank Sandy Mader for the technical support and helpful feedback provided throughout the research process.

Conflicts of Interest: The authors declare no conflict of interest.

Abbreviations

The following abbreviations are used in this manuscript:

AD	automatic differentiation
ADODG	Aerodynamic Design Optimization Discussion Group
AIAA	American Institute of Aeronautics and Astronautics
ASO	aerodynamic shape optimization

CFD	computational fluid dynamics
CRM	Common Research Model
DNS	direct numerical simulation
DPW	Drag Prediction Workshop
FEA	finite element analysis
FFD	free-form deformation
KKT	Karush–Kuhn–Tucker
KS	Kreisselmeier–Steinhauser
LES	large eddy simulation
MAC	mean aerodynamic chord
MACH	MDO of aircraft configurations with high fidelity
MDAO	multidisciplinary analysis and optimization
MDO	multidisciplinary design optimization
MITC	mixed interpolation of tensorial components
OML	outer mold line
RAE	Royal Aircraft Establishment
RANS	Reynolds-averaged Navier–Stokes
TACS	Toolkit for the Analysis of Composite Structures
uCRM	undeflected Common Research Model
XDSM	extended design structure matrix

References

1. Kuchemann, D. *The Aerodynamic Design of Aircraft: A Detailed Introduction to the Current Aerodynamic Knowledge and Practical Guide to the Solution of Aircraft Design Problems*; Pergamon Press: Oxford, UK, 1978.
2. Haftka, R.T. Optimization of Flexible Wing Structures Subject to Strength and Induced Drag Constraints. *AIAA J.* **1977**, *15*, 1101–1106. doi:10.2514/3.7400. [[CrossRef](#)]
3. Chittick, I.R.; Martins, J.R.R.A. Aero-Structural Optimization Using Adjoint Coupled Post-Optimality Sensitivities. *Struct. Multidiscip. Optim.* **2008**, *36*, 59–70. doi:10.1007/s00158-007-0200-9. [[CrossRef](#)]
4. Jansen, P.; Perez, R.E.; Martins, J.R.R.A. Aerostructural Optimization of Nonplanar Lifting Surfaces. *J. Aircr.* **2010**, *47*, 1491–1503. doi:10.2514/1.44727. [[CrossRef](#)]
5. Jasa, J.P.; Hwang, J.T.; Martins, J.R.R.A. Open-source coupled aerostructural optimization using Python. *Struct. Multidiscip. Optim.* **2018**, *57*, 1815–1827. doi:10.1007/s00158-018-1912-8. [[CrossRef](#)]
6. Gray, J.S.; Hwang, J.T.; Martins, J.R.R.A.; Moore, K.T.; Naylor, B.A. OpenMDAO: An open-source framework for multidisciplinary design, analysis, and optimization. *Struct. Multidiscip. Optim.* **2019**, *59*, 1075–1104. doi:10.1007/s00158-019-02211-z. [[CrossRef](#)]
7. Yu, Y.; Lyu, Z.; Xu, Z.; Martins, J.R.R.A. On the Influence of Optimization Algorithm and Starting Design on Wing Aerodynamic Shape Optimization. *Aerosp. Sci. Technol.* **2018**, *75*, 183–199. doi:10.1016/j.ast.2018.01.016. [[CrossRef](#)]
8. Martins, J.R.R.A.; Hwang, J.T. Review and Unification of Methods for Computing Derivatives of Multidisciplinary Computational Models. *AIAA J.* **2013**, *51*, 2582–2599. doi:10.2514/1.J052184. [[CrossRef](#)]
9. Martins, J.R.R.A.; Sturdza, P.; Alonso, J.J. The Complex-Step Derivative Approximation. *ACM Trans. Math. Softw.* **2003**, *29*, 245–262. doi:10.1145/838250.838251. [[CrossRef](#)]
10. Jameson, A. Aerodynamic Design via Control Theory. *J. Sci. Comput.* **1988**, *3*, 233–260. doi:10.1007/BF01061285. [[CrossRef](#)]
11. Kenway, G.K.W.; Mader, C.A.; He, P.; Martins, J.R.R.A. Effective Adjoint Approaches for Computational Fluid Dynamics. *Prog. Aerosp. Sci.* **2019**, *110*, 100542. doi:10.1016/j.paerosci.2019.05.002. [[CrossRef](#)]
12. Martins, J.R.R.A.; Alonso, J.J.; Reuther, J.J. A Coupled-Adjoint Sensitivity Analysis Method for High-Fidelity Aero-Structural Design. *Optim. Eng.* **2005**, *6*, 33–62. doi:10.1023/B:OPTE.0000048536.47956.62. [[CrossRef](#)]
13. Martins, J.R.R.A.; Alonso, J.J.; Reuther, J.J. High-Fidelity Aerostructural Design Optimization of a Supersonic Business Jet. *J. Aircr.* **2004**, *41*, 523–530. doi:10.2514/1.11478. [[CrossRef](#)]
14. Khosravi, S.; Zingg, D.W. Aerostructural Perspective on Winglets. *J. Aircr.* **2017**, *54*. doi:10.2514/1.C033914. [[CrossRef](#)]
15. Khosravi, S.; Zingg, D.W. Aerostructural Optimization of Drooped Wings. *J. Aircr.* **2018**, *55*, 1261–1268. [[CrossRef](#)]

16. Kenway, G.K.W.; Kennedy, G.J.; Martins, J.R.R.A. Scalable Parallel Approach for High-Fidelity Steady-State Aeroelastic Analysis and Derivative Computations. *AIAA J.* **2014**, *52*, 935–951. doi:10.2514/1.J052255. [[CrossRef](#)]
17. Burdette, D.A.; Martins, J.R.R.A. Impact of Morphing Trailing Edge on Mission Performance for the Common Research Model. *J. Aircr.* **2019**, *56*, 369–384. doi:10.2514/1.C034967. [[CrossRef](#)]
18. Brooks, T.R.; Kenway, G.K.W.; Martins, J.R.R.A. Benchmark Aerostructural Models for the Study of Transonic Aircraft Wings. *AIAA J.* **2018**, *56*, 2840–2855. doi:10.2514/1.J056603. [[CrossRef](#)]
19. Brooks, T.R.; Martins, J.R.R.A.; Kennedy, G.J. Aerostructural Trade-offs for Tow-steered Composite Wings. *J. Aircr.* **2020**. doi:10.2514/1.C035699. [[CrossRef](#)]
20. Brooks, T.R.; Martins, J.R.R.A.; Kennedy, G.J. High-fidelity Aerostructural Optimization of Tow-steered Composite Wings. *J. Fluids Struct.* **2019**, *88*, 122–147. doi:10.1016/j.jfluidstructs.2019.04.005. [[CrossRef](#)]
21. Lyu, Z.; Kenway, G.K.W.; Martins, J.R.R.A. Aerodynamic Shape Optimization Investigations of the Common Research Model Wing Benchmark. *AIAA J.* **2015**, *53*, 968–985. doi:10.2514/1.J053318. [[CrossRef](#)]
22. Coder, J.G.; Pulliam, T.H.; Hue, D.; Kenway, G.K.W.; Sclafani, A.J. Contributions to the 6th AIAA CFD Drag Prediction Workshop Using Structured Grid Methods. In *AIAA SciTech Forum*; American Institute of Aeronautics and Astronautics: Reston, VA, USA, 2017. doi:10.2514/6.2017-0960. [[CrossRef](#)]
23. Chernukhin, O.; Zingg, D.W. Multimodality and Global Optimization in Aerodynamic Design. *AIAA J.* **2013**, *51*, 1342–1354. doi:10.2514/1.j051835. [[CrossRef](#)]
24. Koo, D.; Zingg, D.W. Investigation into Aerodynamic Shape Optimization of Planar and Nonplanar Wings. *AIAA J.* **2018**, *56*, 250–263. doi:10.2514/1.J055978. [[CrossRef](#)]
25. Osusky, L.; Buckley, H.; Reist, T.; Zingg, D.W. Drag Minimization Based on the Navier—Stokes Equations Using a Newton—Krylov Approach. *AIAA J.* **2015**, *53*, 1555–1577. doi:10.2514/1.J053457. [[CrossRef](#)]
26. Bons, N.P.; He, X.; Mader, C.A.; Martins, J.R.R.A. Multimodality in Aerodynamic Wing Design Optimization. *AIAA J.* **2019**, *57*, 1004–1018. doi:10.2514/1.J057294. [[CrossRef](#)]
27. Streuber, G.M.; Zingg, D.W. Investigation of multimodality in aerodynamic shape optimization based on the Reynolds-averaged Navier–Stokes equations. In Proceedings of the 18th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference, Denver, CO, USA, 5–9 June 2017.
28. Streuber, G.M.; Zingg, D.W. A Parametric Study of Multimodality in Aerodynamic Shape Optimization of Wings. In *AIAA Aviation Forum*; AIAA: Reston, VA, USA, 2018.
29. Drela, M. *Frontiers of Computational Fluid Dynamics*; Chapter Pros and Cons of Airfoil Optimization; World Scientific: Singapore, 1998; pp. 363–381. doi:10.1142/9789812815774_0019. [[CrossRef](#)]
30. Kenway, G.K.W.; Martins, J.R.R.A. Multipoint Aerodynamic Shape Optimization Investigations of the Common Research Model Wing. *AIAA J.* **2016**, *54*, 113–128. doi:10.2514/1.J054154. [[CrossRef](#)]
31. Reuther, J.J.; Jameson, A.; Alonso, J.J.; Rimlinger, M.J.; Saunders, D. Constrained Multipoint Aerodynamic Shape Optimization Using an Adjoint Formulation and Parallel Computers, Part 2. *J. Aircr.* **1999**, *36*, 61–74. [[CrossRef](#)]
32. Cliff, S.E.; Reuther, J.J.; Saunders, D.A.; Hicks, R.M. Single-Point and Multipoint Aerodynamic Shape Optimization of High-Speed Civil Transport. *J. Aircr.* **2001**, *38*, 997–1005. [[CrossRef](#)]
33. Nemec, M.; Zingg, D.W.; Pulliam, T.H. Multipoint and Multi-Objective Aerodynamic Shape Optimization. *AIAA J.* **2004**, *42*, 1057–1065. [[CrossRef](#)]
34. Gallard, F.; Meaux, M.; Montagnac, M.; Mohammadi, B. Aerodynamic aircraft design for mission performance by multipoint optimization. In Proceedings of the 21st AIAA Computational Fluid Dynamics Conference, San Diego, CA, USA, 24–27 June 2013; American Institute of Aeronautics and Astronautics: Reston, VA, USA, 2013. doi:10.2514/6.2013-2582. [[CrossRef](#)]
35. Mangano, M.; Martins, J.R.R.A. Multipoint Aerodynamic Shape Optimization for Subsonic and Supersonic Regimes. *J. Aircr.* **2020**. doi:10.2514/1.C035826. [[CrossRef](#)]
36. Kenway, G.K.W.; Martins, J.R.R.A. Multipoint High-Fidelity Aerostructural Optimization of a Transport Aircraft Configuration. *J. Aircr.* **2014**, *51*, 144–160. doi:10.2514/1.C032150. [[CrossRef](#)]
37. Liem, R.P.; Kenway, G.K.W.; Martins, J.R.R.A. Multimission Aircraft Fuel Burn Minimization via Multipoint Aerostructural Optimization. *AIAA J.* **2015**, *53*, 104–122. doi:10.2514/1.J052940. [[CrossRef](#)]
38. Wakayama, S.; Kroo, I. Subsonic Wing Planform Design Using Multidisciplinary Optimization. *J. Aircr.* **1995**, *32*, 746–753. [[CrossRef](#)]

39. Ning, A.; Kroo, I. Multidisciplinary Considerations in the Design of Wings and Wing Tip Devices. *J. Aircr.* **2010**, *47*, 534–543. doi:10.2514/1.41833. [[CrossRef](#)]
40. Buckley, H.P.; Zhou, B.Y.; Zingg, D.W. Airfoil Optimization Using Practical Aerodynamic Design Requirements. *J. Aircr.* **2010**, *47*, 1707–1719. doi:10.2514/1.C000256. [[CrossRef](#)]
41. Vassberg, J.C.; DeHaan, M.A.; Rivers, S.M.; Wahls, R.A. Development of a Common Research Model for Applied CFD Validation Studies. In Proceedings of the 26th AIAA Applied Aerodynamics Conference, Honolulu, HI, USA, 18–21 August 2008. doi:10.2514/6.2008-6919. [[CrossRef](#)]
42. Lambe, A.B.; Martins, J.R.R.A. Extensions to the Design Structure Matrix for the Description of Multidisciplinary Design, Analysis, and Optimization Processes. *Struct. Multidiscip. Optim.* **2012**, *46*, 273–284. doi:10.1007/s00158-012-0763-y. [[CrossRef](#)]
43. Gill, P.E.; Murray, W.; Saunders, M.A. *User's Guide for SNOPT Version 7: Software for Large-Scale Nonlinear Programming*; Technical Report; Systems Optimization Laboratory, Stanford University: Stanford, CA, USA, 2007.
44. Sederberg, T.W.; Parry, S.R. Free-form Deformation of Solid Geometric Models. *SIGGRAPH Comput. Graph.* **1986**, *20*, 151–160. doi:10.1145/15886.15903. [[CrossRef](#)]
45. Kenway, G.K.; Kennedy, G.J.; Martins, J.R.R.A. A CAD-Free Approach to High-Fidelity Aerostructural Optimization. In Proceedings of the 13th AIAA/ISSMO Multidisciplinary Analysis Optimization Conference, Fort Worth, TX, USA, 13–15 September 2010; Number AIAA 2010-9231. doi:10.2514/6.2010-9231. [[CrossRef](#)]
46. Yildirim, A.; Kenway, G.K.W.; Mader, C.A.; Martins, J.R.R.A. A Jacobian-free approximate Newton–Krylov startup strategy for RANS simulations. *J. Comput. Phys.* **2019**, *397*, 108741. doi:10.1016/j.jcp.2019.06.018. [[CrossRef](#)]
47. Kennedy, G.J.; Martins, J.R.R.A. A Parallel Finite-Element Framework for Large-Scale Gradient-Based Design Optimization of High-Performance Structures. *Finite Elem. Anal. Des.* **2014**, *87*, 56–73. doi:10.1016/j.finel.2014.04.011. [[CrossRef](#)]
48. Brown, S.A. Displacement Extrapolation for CFD+CSM Aeroelastic Analysis. In Proceedings of the 35th AIAA Aerospace Sciences Meeting, Reno, NV, USA, 6–9 January 1997; AIAA: Reston, VA, USA, 1997.
49. Kennedy, G.J.; Martins, J.R.R.A. A Comparison of Metallic and Composite Aircraft Wings Using Aerostructural Design Optimization. In Proceedings of the 14th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference, Indianapolis, IN, USA, 17–19 September 2012; AIAA-2012-5475. doi:10.2514/6.2012-5475. [[CrossRef](#)]
50. U.S. Standard Atmosphere. *NASA Technical Memorandum NASA-TM-X-74335, NOAA-S/T-76-1562*; NASA: Washington, DC, USA, 1976.
51. Lambe, A.B.; Martins, J.R.R.A.; Kennedy, G.J. An Evaluation of Constraint Aggregation Strategies for Wing Box Mass Minimization. *Struct. Multidiscip. Optim.* **2017**, *55*, 257–277. doi:10.1007/s00158-016-1495-1. [[CrossRef](#)]
52. Kenway, G.K.W.; Martins, J.R.R.A. Buffet Onset Constraint Formulation for Aerodynamic Shape Optimization. *AIAA J.* **2017**, *55*, 1930–1947. doi:10.2514/1.J055172. [[CrossRef](#)]
53. Toal, D.J.J.; Keane, A.J. Efficient Multipoint Aerodynamic Design Optimization Via Cokriging. *J. Aircr.* **2011**, *48*, 1685–1695. doi:10.2514/1.c031342. [[CrossRef](#)]

