

A Comparison of Metallic and Composite Aircraft Wings Using Aerostructural Design Optimization

Graeme J. Kennedy,*

University of Toronto Institute for Aerospace Studies, Toronto, ON, Canada

Joaquim R. R. A. Martins†

University of Michigan, Department of Aerospace Engineering, Ann Arbor, MI, USA

In this paper we examine the design of metallic and composite aircraft wings in order to assess how the use of composites modifies the trade-off between structural weight and drag. In order to perform this assessment, we use a gradient-based aerostructural design optimization framework that combines a high-fidelity finite-element structural model that includes panel-level design variables with a medium fidelity aerodynamic panel method with profile and compressibility drag corrections. In order to examine the effect of the choice of the objective, we obtain a Pareto front of designs by minimizing a weighted combination of the mission fuel burn and take-off gross-weight of the aircraft over a multi-segment mission profile. The structural model includes both strength and buckling constraints and includes a detailed laminate parametrization that is used to obtain the optimal lamination stacking sequence and impose manufacturing requirements for composites including matrix-cracking and minimum ply-content constraints. We show that the composite wing designs are between 34% and 40% lighter than the equivalent metallic wings. Due to this large structural weight savings, the composite aircraft designs exhibit a fuel burn savings of between 5% and 8% and a take-off gross-weight savings of between 6% and 11%.

I. Introduction

In preliminary aircraft design studies, aircraft weight estimates are often obtained based on simplified models that are calibrated with historical data. These types of weight models are often correlations, sometimes enhanced with physical reasoning, that are computationally inexpensive to use and usually provide sufficiently accurate weight predictions within a limited design space. However, such weight prediction methods have several drawbacks. First, when using these types of methods, it may be difficult to assess the relative benefits of novel structural technologies that may reduce the structural wing weight compared to conventional designs. Second, these types of models may not provide sufficient accuracy for new materials, new structural arrangements, unconventional aircraft configurations or novel wingtip devices. In this paper, we use high-fidelity structural analysis methods to predict the structural wing weight and compare results for metallic and composite wing constructions. In addition, we examine the wing design for both composite and metallic constructions as a function of a weighted objective of fuel burn and takeoff gross-weight. This enables us to make a comparison between metallic and composite constructions over a range of objectives that place emphasis on either drag or structural weight.

It has long been understood that there is a fundamental trade-off between drag and aircraft weight. This trade-off has been examined by many authors. For instance, Jones [1] presented an analysis of wings with minimum induced drag for fixed lift and root bending moment. Jones found that a 15% reduction in the induced drag could be obtained by increasing the span 15% for a fixed root bending moment. Later, Jones and Lasinski [2] presented an analysis of nonplanar lifting surfaces using an integrated bending moment constraint. They found that winglets and wing tip extensions provided approximately equal reduction in the induced drag for a given integrated bending moment constraint. More recently, Ning and Kroo [3] performed an analysis and optimization of wings with various wing tip devices. They used a weight model that included a historical weight correlation and an integrated bending moment over wing thickness to predict relative changes in weight of different designs.

Other studies have used structural analysis techniques to obtain the flying shape of the wing, and to size a portion of the aircraft structure and thus predict a partial wing-weight. Haftka [4], compared the trade-off between structural weight and induced drag for both composite and isotropic wings of a fighter aircraft. Haftka obtained the displaced shape through an iterative procedure and used stress constraints to size the aircraft wing skins. More recently, Jansen *et al.* [5] presented optimizations of various nonplanar configurations using a gradient-free optimization method. They

*Ph.D. Candidate, AIAA Student Member

†Associate Professor, AIAA Senior Member

used a calibrated lifting line method and obtained the displaced shape of each design. This paper a bit more. They obtained a box wing, a wing with winglets, a C-wing, and a wing with raked wingtips, depending on the design formulation, whether structures was considered, and which types of drag were included in their drag model.

Many authors have developed techniques for sizing either isotropic or composite wing-box structures for strength, buckling, and manufacturing constraints. Liu *et al.* [6] performed a global-local optimization of a composite wing with unstiffened panels. The formulation included a global optimization problem to minimize structural weight, subject to failure constraints, while the local optimization problem included a stacking sequence design for maximum buckling load. A response surface model was used within the formulation to construct an approximation of the locally optimum designs as a function of the global design variables. Later, Liu and Haftka [7] performed a single-level optimization using lamination parameters and showed that their results were identical to the two-level formulation of Liu *et al.* [6]. In order to place the two-level optimization approach on a more rigorous foundation, Haftka and Watson [8] developed a decomposition theory for a class of quasi-separable optimization problems. Liu *et al.* [9] applied this quasi-separable approach to a beam-frame weight-minimization problem.

High-fidelity aerostructural analysis and design optimization methods have also been applied to aircraft design problems. Martins *et al.* [10] used an Euler CFD solver coupled to a linear finite-element solver to design a business jet using a linearized range objective with an adjoint-based gradient evaluation method that enabled optimization with respect to hundreds of design variables that included aerodynamic shape and structural sizing variables [11]. More recently, Kenway *et al.* [12], used an Euler CFD solver coupled to a parallel finite-element solver to design a transport aircraft for minimum fuel burn and minimum takeoff gross-weight.

In this paper we combine a high-fidelity finite-element structural model that includes panel-level design variables with a medium fidelity aerodynamic panel method. Our goal is to assess the trade-off between structural weight and drag and to determine how the use of composite materials modifies these trade-offs. In this study, the use of gradient-based optimization techniques is essential, since the structural panel-level parametrization requires thousands of design variables. Our framework is based on the high-fidelity aerostructural optimization approach developed by Martins *et al.* [10] and the aerostructural implementation of Kennedy and Martins [13], and is part of the framework presented by Kenway *et al.* [12].

The remainder of this paper is organized as follows: In Section II we briefly describe the aerostructural analysis and optimization framework used in this work. In Section III we discuss the details of key components of the analysis that are of significance to this study. In Section IV we outline the aerostructural design optimization problem for both the composite and metallic wings. In Section V we present the aerostructural design cases and compare the composite and metallic designs.

II. Aerostructural analysis and gradient-evaluation methods

In this work, all constraint functions and the objective are evaluated using converged aerostructural solutions, and therefore, an efficient aerostructural solution method is required to obtain results within reasonable computational time. Furthermore, due to the large dimensionality of the design space and the computational cost of the analysis, we exclusively use gradient-based design optimization methods with an efficient adjoint-based gradient evaluation method. The following section outlines the aerostructural analysis and adjoint-based gradient evaluation techniques used within this work. Additional details about the solution methods and the efficiency of the approach can be found in Kennedy and Martins [13].

A. Aerodynamic analysis

The aerodynamic analysis is performed using TriPan, an unstructured, three-dimensional parallel panel code for calculating the aerodynamic forces, moments and pressures for inviscid, incompressible, external lifting flows using the Prandtl–Glauert equation [13]. TriPan uses constant first-order source and doublet singularity elements distributed over the lifting surface and doublet elements distributed over the wake [14]. The source strengths are determined based on the onset flow conditions while the boundary conditions for the doublet strengths constitute a dense linear system, represented here by

$$\mathbf{R}_A(\mathbf{w}, \mathbf{u}) = 0, \quad (1)$$

where \mathbf{u} and \mathbf{w} are vectors of the structural and aerodynamic state variables, respectively. The linear system represented by Equation (1) is solved in parallel using PETSc [15, 16]. A dense matrix format is used for the matrix-vector products, while a sparse approximate-Jacobian is used to form a incomplete LU (ILU) preconditioner. The linear system is solved using the Krylov subspace method GMRES.

B. Load transfer

The load and displacement transfer scheme follows the work of Brown [17]. The displacements from the structures are extrapolated to the aerodynamic nodes using rigid links. These rigid links are formed by locating the closest point on the structural surface to each of the aerodynamic nodes. The structural surface is determined by interpolating between structural nodes using the finite-element shape functions. The displacements \mathbf{u}_S and rotations $\boldsymbol{\theta}_S$ on the structural surface, and the rigid links \mathbf{r} are used to determine the displacements of the aerodynamic nodes \mathbf{u}_A as follows:

$$\mathbf{u}_A = \mathbf{u}_S + \boldsymbol{\theta}_S \times \mathbf{r}. \quad (2)$$

Note that this formula uses a small angle approximation. Equation (2) can be used, in conjunction with the method of virtual work, to form the consistent force vector for the aerodynamic forces at the structural nodes. More details of the approach are outlined in Kennedy and Martins [13].

C. Structural analysis

The structural analysis is performed using the Toolkit for the Analysis of Composite Structures (TACS), a parallel, finite-element code developed by the authors, designed specifically for the analysis of stiffened, thin-walled, composite structures using either linear or geometrically nonlinear strain relationships [13]. We typically use higher-order finite-elements as we have found that these provided better stress prediction capability. The residuals of the structural governing equations are

$$\mathbf{R}_S(\mathbf{w}, \mathbf{u}) = \mathbf{S}_c(\mathbf{u}) - \mathbf{F}(\mathbf{w}, \mathbf{u}), \quad (3)$$

where \mathbf{u} is a vector of displacements and rotations, \mathbf{S}_c are the residuals due to conservative forces and internal strain energy and \mathbf{F} are the follower forces due to aerodynamic loads.

The Jacobian of the structural residuals involves two terms: the tangent stiffness matrix $\mathbf{K} = \partial \mathbf{S}_c / \partial \mathbf{u}$ and the derivative of the consistent force vector with respect to the structural displacements. This results in the following expression for the Jacobian of the structural residuals:

$$\frac{\partial \mathbf{R}_S}{\partial \mathbf{u}} = \mathbf{K} - \frac{\partial \mathbf{F}}{\partial \mathbf{u}}. \quad (4)$$

While the matrices involved in structural problems are typically symmetric, the term $\partial \mathbf{F} / \partial \mathbf{u}$ is non-symmetric due to the non-conservative nature of the aerodynamic forces. These non-symmetric matrices require different solution algorithms than those typically employed in structural finite-element codes. We use GMRES [18] to solve the non-symmetric, linear systems involving the matrix in Equation (4).

D. Approximate Newton–Krylov method

The aerostructural residuals are the concatenation of the aerodynamic and structural residuals, represented by:

$$\mathbf{R}(\mathbf{q}, \mathbf{x}) = \begin{bmatrix} \mathbf{R}_A(\mathbf{w}, \mathbf{u}, \mathbf{x}) \\ \mathbf{R}_S(\mathbf{w}, \mathbf{u}, \mathbf{x}) \end{bmatrix} = 0, \quad (5)$$

where \mathbf{R}_A and \mathbf{R}_S are the aerodynamic and structural residuals, \mathbf{w} and \mathbf{u} are the aerodynamic and structural state variables, \mathbf{q} is the full set of aerostructural state variables $\mathbf{q}^T = [\mathbf{w}^T, \mathbf{u}^T]$, and \mathbf{x} is a vector of design variables.

Newton's method applied to Equation (5) results in the following linear system of equations for the update $\Delta \mathbf{q}^{(n)}$,

$$\frac{\partial \mathbf{R}}{\partial \mathbf{q}} \Delta \mathbf{q}^{(n)} = -\mathbf{R}(\mathbf{q}^{(n)}). \quad (6)$$

In the Newton–Krylov approach, the linearized system (6) is solved inexactly using a Krylov subspace method. Here we use a preconditioner based on generic discipline-level preconditioners that ignores the off-diagonal coupling terms. The off-diagonal matrix-vector products are computed using a product-rule implementation that is discussed in further detail in Kennedy and Martins [13].

E. Adjoint-based gradient-evaluation

Efficient gradient-based optimization requires the accurate and efficient evaluation of gradients of functions of interest. In the aerostructural optimization problem there are typically far fewer objective and constraint functions than there are design variables. Therefore, compute the required gradients using the adjoint method is advantageous. We have

developed an aerostructural adjoint that is based entirely on analytic derivatives without the use of finite-difference computations. The coupled aerostructural adjoint equations can be written in the following form:

$$\frac{\partial \mathbf{R}^T}{\partial \mathbf{q}} \boldsymbol{\psi} = \frac{\partial f^T}{\partial \mathbf{q}}, \quad (7)$$

where $\boldsymbol{\psi}$ is the adjoint vector and $f(\mathbf{q}, \mathbf{x})$ is either an aerodynamic or structural function of interest. Once the adjoint vector has been determined using Equation (7), the total derivative is determined using the additional computation:

$$\nabla_{\mathbf{x}} f = \frac{\partial f}{\partial \mathbf{x}} - \boldsymbol{\psi}^T \frac{\partial \mathbf{R}}{\partial \mathbf{x}}. \quad (8)$$

We use a Krylov method to solve the linear coupled aerostructural adjoint equations (7) in an analogous manner to the Krylov method applied to the linearized Newton system. In the Krylov approach, the matrix-vector products are computed using the exact Jacobian-transpose of the coupled aerostructural system. One iteration of a transpose block Jacobi iteration is used as the preconditioner.

III. Design problem components

The aerostructural analysis enables the prediction of the displacements, stresses and aerodynamic forces on the displaced, aerostructural system. Within the context of aerostructural design optimization, these results must be combined into a consistent objective and constraint formulation that reflects the most important aspects of the aircraft design problem. In this section, we describe in detail the analysis used to predict the drag and maximum lift of the aircraft, as well as the composite parametrization method used for the structural wing design and the panel-level analysis used to size the structure for buckling.

A. Drag analysis

TriPan is an inviscid code that can be used to accurately compute the induced drag, but is unable to account for wave and profile drag contributions that are important considerations in wing design. Here, we model these additional drag contributions using semi-empirical methods.

The profile drag is computed based on a quadratic model of the sectional drag coefficient:

$$c_{dp} = c_{d0} + c_{d2} c_l^2, \quad (9)$$

where c_{dp} is the profile drag. The coefficient c_{d0} is based on the skin friction estimate

$$c_{d0} = F_c c_f,$$

where F_c is a profile drag form-factor, and c_f is the turbulent skin-friction coefficient determined using the van Driest II method [19]. The form-factor, F_c , is computed using the thickness to chord ratio, t/c , as follows:

$$F_c = 1 + 2.7 \left(\frac{t}{c} \right) + 100 \left(\frac{t}{c} \right)^4.$$

Finally, based on the method presented by Wakayama and Kroo [20], the quadratic coefficient in Equation (9), is computed based on the expression:

$$c_{d2} = \frac{0.38}{\cos^2 \Lambda} c_{d0}.$$

The compressibility drag is computed based on a crest-critical Mach number computed using the Korn equation,

$$M_{\text{crit}} = \frac{\kappa_A}{\cos \Lambda} - \frac{t/c}{\cos^2 \Lambda} - \frac{c_l}{10 \cos^3 \Lambda} - \left(\frac{0.1}{80} \right)^{1/3}, \quad (10)$$

where κ_A is a technology factor that we set to $\kappa_A = 0.95$, which is suitable for supercritical airfoil sections commonly used on transport aircraft. The sectional contribution to the compressibility drag is then computed using

$$c_{dc} = 20(M - M_{\text{crit}})^4 \quad (11)$$

for $M > M_{\text{crit}}$.

Figure 1 shows the variation of C_D , including induced, profile and compressibility drag, with increasing Mach number for fixed C_L , for the reference configuration described in Section II. The decreasing C_D values for increasing Mach number are due to Reynolds number effects, while the drag divergence behavior due to wave drag can clearly be seen above $M = 0.8$.




Figure 1: C_D for increasing Mach number at constant C_L

B. Maximum lift prediction

The maximum lift capability of a wing is an important factor in wing design. In this work, we employ a critical sectional analysis approach in which the maximum lift capability of a wing is determined when any one airfoil section reaches its maximum lift capability. In order to assess the maximum lift for each section, we use the Valarezo criterion [21]. The Valarezo criterion is based on the absolute value of the difference between the peak leading edge pressure coefficient and the trailing edge pressure coefficient, which we label ΔC_p . Valarezo and Chin [21] compared a large collection of C_p data for sections at maximum lift and found that the ΔC_p values for all sections collapsed to an allowable ΔC_p that depends on Mach and Reynolds number, but is independent of the airfoil shape or on the arrangement of any high-lift devices. Using a large collection of sectional pressure data, they were able to predict, within reasonable bounds, the maximum lift capability of several aircraft wings, with and without high-lift devices [21].




Figure 2: The sectional C_p and ΔC_p values for an untwisted wing.

We have implemented the Valarezo criterion within TriPan to predict the maximum lift capability for clean aircraft wings, without high-lift devices. Figure 2 shows the C_p distribution at maximum lift for an untwisted wing that is a linear loft of RAE2822 airfoil sections, at a Mach number of $M = 0.25$ and an altitude of 20 000 ft. This wing is described in more detail in Section II. The C_p distribution shown in Figure 2 corresponds to a $C_{L_{\max}} = 1.03$, where the critical section occurs at a span-wise station of approximately $\eta = 0.9$. The largest allowable ΔC_p is uniform value of 14 independent of chord-wise Reynolds number. However at slower speeds and lower Reynolds numbers, the allowable ΔC_p would vary span-wise.

One of the advantages of the Valarezo constraint is that it is capable of predicting the variation of $C_{L_{\max}}$ with geometry changes. In the present study, we are primarily concerned with design variables that exhibit strong coupling between both aerodynamics and structures. One of the most important variables in this category is the thickness-to-chord ratio. Figure 3 shows the variation of c_l distribution at $C_{L_{\max}}$, and the $C_{L_{\max}}$ variation with t/c ratios between




Figure 3: C_{Lmax} as a function of t/c predicted using the Valarezo stall condition.

9% and 15%. The present analysis predicts a linear relationship between decreasing t/c ratio and decreasing C_{Lmax} , whereas typically a nonlinear relationship would exist, peaking between 12% and 20% t/c ratio. This discrepancy is due to the lack of a boundary-layer correction within TriPan. However, this linear variation is sufficient for a preliminary estimate and captures the overall trend of decreasing maximum lift capacity with decreasing t/c ratio.

In order to be useful within the context of an optimization problem, the Valarezo criterion must be formulated as a constraint. To formulate this constraint we use the ΔC_p margin, defined as follows:

$$\Delta C_{pmargin}(\eta) = \Delta C_{pallow} - \Delta C_p. \quad (12)$$

Using the Valarezo criterion the minimum ΔC_p margin is zero at C_{Lmax} , such that

$$\min \Delta C_{pmargin}(\eta) = 0.$$

For the purposes of optimization, we examine the ΔC_p margin at a series of span locations and obtain an approximate minimum value using the Kreisselmeier–Steinhauser (KS) aggregation function [22]. This aggregated constraint can be written as follows:

$$KS_{\min}(\Delta C_{pmargin}(\eta_k), \rho) = 0 \quad (13)$$

where η_k are the span-wise locations and $KS_{\min}(\cdot, \rho)$ is the KS aggregation function with parameter ρ , where we typically use a value of $\rho = 20$.

C. Composite structural parametrization

In this work, we use a composite parametrization that includes the full through-thickness lamination sequence. In order to obtain a feasible lamination sequence that includes common manufacturing constraints, we use an extension of the technique developed by Kennedy and Martins [23] that can be used to construct a lamination sequences from a set of discrete allowable ply angles. Like previous authors, we restrict the allowable ply angles to the set of angles $\Theta = \{-45^\circ, 0^\circ, 45^\circ, 90^\circ\}$. The laminate parametrization approach presented in Kennedy and Martins [23], is a continuous relaxation of the laminate parametrization first developed by Le Riche and Haftka [24], coupled with an exact ℓ_1 penalization to enable gradient-based design optimization. In the present approach, the lamination sequence is expressed in terms of continuous ply-identity design variables over the interval $x_{ij} \in [0, 1]$. An active ply variable, $x_{ij} = 1$, indicates that the j^{th} discrete ply angle from the set Θ is active in the i^{th} ply. Using the ply-identity variables, the stiffness matrices, \mathbf{A} and \mathbf{D} , for a symmetric laminate can be expressed as follows:

$$\begin{aligned} \mathbf{A} &= 2 \frac{t}{N} \sum_{i=1}^N x_{ij} \sum_{j=1}^4 \bar{\mathbf{Q}}(\theta_j) \\ \mathbf{D} &= \frac{2}{3} \left(\frac{t}{N} \right)^3 \sum_{i=1}^N (i^3 - (i-1)^3) \sum_{j=1}^4 x_{ij} \bar{\mathbf{Q}}(\theta_j) \end{aligned} \quad (14)$$

where t is the total thickness of the laminate, and $\bar{\mathbf{Q}}(\theta)$ is the lamina thickness in the laminate reference frame.

In order to obtain a feasible design, constraints must be imposed to ensure that only one ply-identity variable is active at a time in a single layer. This condition can be imposed by enforcing two simultaneous constraints: that the design lies on the unit hyper-plane

$$\sum_{j=1}^4 x_{ij} = 1, \quad i = 1, \dots, N, \quad (15)$$

and that the design lies on the unit sphere

$$\sum_{j=1}^M x_{ij}^2 = 1, \quad i = 1, \dots, N. \quad (16)$$

At feasible points only one ply-identity variable is active in a given layer while all other ply-identity variables are zero. For ease of presentation we collect the linear constraints in Equation (15) for all N plies into the matrix expression:

$$\mathbf{A}_w \mathbf{x} = \mathbf{e}, \quad (17)$$

where $\mathbf{e} \in \mathbb{R}^N$ is a vector of unit entries. We also collect the nonlinear spherical constraints from Equation (16) for all N plies into the vector:

$$\mathbf{c}_s(\mathbf{x}) = \mathbf{e}. \quad (18)$$

One issue with this formulation is that the second constraint introduces many local minima into the design problem. To overcome this problem, we use an ℓ_1 penalization approach such that the original objective $f(\mathbf{x})$ is replaced by a modified objective $f(\mathbf{x}) + \gamma \|\mathbf{c}_s(\mathbf{x}) - \mathbf{e}\|_1$, where $\|\cdot\|_1$ is the ℓ_1 norm. We solve a sequence of optimization problems, while gradually increasing γ , until the spherical constraints are satisfied, i.e. $\mathbf{c}_s(\mathbf{x}) = \mathbf{e}$. The advantage of the ℓ_1 penalty function is that it is exact such that solutions to the original problem are also solutions to the penalized problem.

As shown in Kennedy and Martins [23], when the linear constraints (15) are satisfied exactly at every iteration, the ℓ_1 norm can be replaced by the following expression:

$$\|\mathbf{c}_s(\mathbf{x}) - \mathbf{e}\|_1 = \mathbf{e}^T (\mathbf{e} - \mathbf{c}_s(\mathbf{x})). \quad (19)$$

This simplification modifies the objective such that it is differentiable and amenable to gradient-based optimization. In the present work we use the optimization code SNOPT [25], through the Python interface supplied by the pyOpt optimization package [26]. SNOPT satisfies the linear constraints at every iteration, enabling the use of Equation (19).

While the laminate parametrization presented by Kennedy and Martins [23] can handle general non-symmetric laminates, here we impose additional constraints on the lamination sequence as proposed by Baker *et al.* Chap. 12 [27] for practical lamination sequences:

1. Each laminate must be balanced, containing equal numbers of $\pm 45^\circ$ plies,
2. Each laminate must contain at least 10% of plies in each of the directions $\Theta = \{-45^\circ, 0^\circ, 45^\circ, 90^\circ\}$,
3. At most four adjacent plies in each laminate can be in the same direction.

We enforce the balanced and 10% ply content constraint not just to the laminate as a whole, but within every sub-sequence of 10 plies starting from the middle layer and proceeding outwards. If fewer than 5 plies are remaining in the last sub-sequence, we remove the ply content constraint. Imposing the ply content constraint in this manner ensures that all ply angles are distributed throughout the laminate.

The balanced laminate constraint is enforced as follows:

$$\sum_{i=10(k-1)+1}^{10k+1} (x_{i1} - x_{i3}) = 0, \quad k = 1, \dots, N/10, \quad (20)$$

where $j = 1$ and $j = 3$ correspond to the -45° and 45° plies respectively, and k indicates the index of the 10-ply laminate sub-sequence. In a similar manner, the 10% ply content constraint is formulated as follows:

$$\sum_{i=10(k-1)+1}^{10k+1} x_{ij} = 1, \quad j = 1, \dots, 4 \quad k = 1, \dots, N/10, \quad (21)$$

such that one of each ply from each member of the set Θ must be active in every sub-sequence of 10 plies.

To avoid lamination sequences with more than four contiguous plies at any given angle, we impose the following constraint:

$$\sum_{i=k}^{k+5} x_{ij} \leq 4, \quad j = 1, \dots, 4 \quad k = 1, \dots, N - 5, \quad (22)$$

and at the symmetry plane, we impose the following additional constraints:

$$\sum_{i=0}^k x_{ij} \leq k, \quad j = 1, \dots, 4 \quad k = 2, \dots, 4. \quad (23)$$

These constraints ensure that over a 5-ply sequence, no more than four plies are active and that at the symmetry plane, no more than two plies may be active.

In order to provide a parametrization scheme that includes a failure criterion, we adopt a conservative failure envelope that includes contributions from all allowable ply angles $\theta_j \in \Theta$. We use the Tsai–Wu failure criterion, which can be written as follows:

$$F(\sigma_l(\theta)) \leq 1, \quad (24)$$

where $\sigma_l(\theta)$ is the laminae stress at an angle θ to the laminate reference frame. To construct the overall laminate failure envelope, we apply the failure criteria (24) at all angles $\theta_j \in \Theta$, at the upper and lower surfaces of the laminate. Instead of applying each of these criteria independently, we aggregate them into a single function using the KS aggregation technique [22]. This provides a conservative failure envelope, but does not account for the variation of ply angles within the laminate. This conservative failure envelope can be written as follows:

$$F_{\text{KS}}^{(i)}(\sigma) = \text{KS}(F(\sigma_l^{(p)}), \rho) \leq 1, \quad (25)$$

where the aggregation takes place over the range $p = 1, \dots, 8$. Here, $\sigma_l^{(2j-1)}$ and $\sigma_l^{(2j)}$ are the laminae stresses at the angle θ_j on the top and bottom surfaces, respectively. Equation (25) provides a conservative failure envelope in the sense that when the laminate stresses are within the envelope, all laminae within the layup are within the failure envelope represented by Equation (24).

The preceding discussion has focused on a fixed-thickness laminate with a variable stacking sequence. However, practical sizing problems require a variable thickness in order to size the structure so that it is not heavier than absolutely necessary. To address this issue, we select an initial thickness distribution that we use to determine the number of plies in each structural component. During the optimization we keep the number of plies in each component fixed, but allow the laminate thickness to change by scaling the ply thicknesses by the same factor in all layers. As a result, in the optimized design there is often a miss-match between the physical ply thickness, and scaled ply thickness. The severity of this miss-match can be assessed by examining the discrepancy between the number of fixed plies selected initially and the number of physical ply thicknesses required for each component. If this discrepancy is too large, the optimization problem can be restarted using the optimized thickness distribution to predict the number of plies in each component. We examine the effect of such a procedure in Section V.

D. Panel-level analysis

The finite-element model of the wing is constructed using a smeared stiffness approach where the stiffness contribution from the stringers is added to the effective stiffness of the skin. Once the solution to the aerostructural system is obtained, the internal loads from the finite-element model are used to evaluate failure and buckling constraints within a series of local panel models. These local panel models are formed by the spar-rib intersections in the wing and include discrete stiffeners. The failure constraints for each panel are evaluated using simple mechanics of materials relationships, while the buckling constraints are evaluated using a buckling-interaction envelope that is calculated using a finite-strip method.

The panel-level design variables are shown in Figure 4. These variables include both panel geometry variables and panel thicknesses. The panel geometry variables include the stiffener pitch, b , the stiffener height, h_s , the base width, w_b , and the flange width, f_w . The thickness variables are the skin thickness, t_s , the stiffener thickness, t_w , and the base thickness t_b . It is important to note that when composite materials are used, the laminate stacking sequence parametrization also modify the stiffness of the different panel components and must be included as panel-level design variables.

For the failure constraints, the strain from the global finite-element model is used in conjunction with strain simple mechanics of materials relationships to obtain the strain at critical points within the section. These points include




Figure 4: The panel dimensions and thicknesses used as design variables within the optimization problem. The dimensions are: b , the stiffener pitch, h_s , the stiffener height, w_b , the base width, and f_w , the flange width. The thicknesses are t_s , the skin thickness, t_w , the stiffener thickness, and $t_b = t_s + 1/2t_w$ the base thickness.

the flange section, stiffener web, skin and base-web intersection. The strain at these points is used as input to the failure criterion (25), or in the metallic case, the von Mises criterion. Finally, the failure envelope from each element is aggregated over the global finite-element model using the KS function. The aggregated KS failure constraint can be written as follows:

$$\text{KS}(F_{\text{KS}}(\boldsymbol{\sigma}), 30), 50) \leq 1, \quad (26)$$

where an aggregation parameter of $\rho = 50$ is used.

The buckling constraints are imposed by constructing a buckling-interaction envelope from the critical loads of an equivalent panel computed using a finite-strip analysis. The critical loads are determined under the assumption that the panels are nearly flat, and so no curvature effects are captured. We assume that the interaction between the longitudinal and shear buckling modes collapses into the following buckling envelope:

$$B(\bar{N}_x, \bar{N}_{xy}) = F_1 \bar{N}_x + F_2 \bar{N}_{xy} + F_{22} \bar{N}_{xy}^2 \leq 1, \quad (27)$$

where \bar{N}_x and \bar{N}_{xy} are the average longitudinal and average shear forces in a local reference frame where the x -direction is aligned with the stringers. In this locally aligned axis, the ribs may not be perpendicular to the stringers and the panels may be stiffened parallelograms. This effect is captured by the finite-strip model and results in unequal positive and negative shear buckling modes. The parameters F_1 , F_2 and F_{22} in Equation (27), are given by:

$$\begin{aligned} F_1 &= -\frac{1}{N_{x,cr}}, \\ F_2 &= \frac{N_{xy,cr}^+ - N_{xy,cr}^-}{N_{xy,cr}^+ N_{xy,cr}^-}, \\ F_{22} &= \frac{1}{N_{xy,cr}^+ N_{xy,cr}^-}, \end{aligned}$$

where, $N_{x,cr}$ is the compressive buckling load of the panel and $N_{xy,cr}^+$ and $N_{xy,cr}^-$ are the positive and negative shear buckling loads. Note that the positive and negative shear buckling loads are not equal when the stringers are not perpendicular to the ribs. The average longitudinal and shear resultants, \bar{N}_x and \bar{N}_{xy} , are obtained for each panel using the following formula:

$$\begin{aligned} \bar{N}_x &= \frac{1}{A} \int N_x \, dA, \\ \bar{N}_{xy} &= \frac{1}{A} \int N_{xy} \, dA, \end{aligned}$$

where A is the area of the local panel. As a result of this formulation, there is only one panel buckling constraint per panel within the model. For typical finite-element problems, there are far fewer buckling constraints than there are stress constraints.

E. Inertial relief

Inertial relief from fuel weight and structural self-weight may produce significant loads on the wing. In the analysis and optimization results presented here, we use both structural self-weight and inertial relief due to fuel in all aerostructural calculations. All inertial relief loads are calculated based on an inertial vector, \mathbf{n} , which is a scalar multiple of the acceleration due to gravity. The structural self-weight is calculated in a straightforward manner using a consistent self-weight load using the finite-element shape functions.

The inertial relief due to fuel is calculated using a set of compatibility variables and consistency constraints that are added to the optimization problem. In this formulation, the mass of the fuel in the wing at an operating point is determined from a mission analysis and is added as a variable within the optimization problem. This fuel mass variable is denoted m_{fuel} . The fuel load is calculated over either the upper or lower wing skin, depending on the direction of the inertial vector \mathbf{n} , using a mass-per-unit area variable ρ_k , where k denotes the k^{th} rib-bay. A series of consistency constraints are used to ensure that the total fuel load is applied to the wing, and that the local load is proportional to the volume of fuel in the present bay. These consistency constraints are

$$A_k \rho_k = m_{\text{fuel}} \frac{V_k}{V_{\text{fuel}}}, \quad (28)$$

where A_k is the area of the skin for the rib-bay, V_k is the volume of the rib-bay and V_{fuel} is the volume over which the fuel is distributed such that, $V_{\text{fuel}} = \sum_k V_k$. The advantage of this approach is that the total mass of the fuel is added to structure regardless of the volume distribution in the wing. Thus, even at intermediate designs, the full fuel load is applied to the wing. However, this constraint does not ensure that there is enough volume within the wing to fit the required mission fuel. In order to ensure that the mission fuel fits within the wing, we also add a volume constraint to the optimization problem such that the maximum range mission fuel will fit within the volume enclosed by the leading and trailing edge spars of the wing.

F. Geometry parametrization

The geometric parametrization of the aerodynamic surfaces and structural surfaces and volumes, including all internal structure, is a key component of aerostructural design optimization. Here, we use a CAD-free approach to manipulate the underlying discipline-level meshes in a continuous and differentiable manner that is well-suited for aerostructural design optimization problems [28]. The geometric parametrization uses a free-form deformation (FFD) [29] approach that defines a modification or deformation of the initial geometry. In the FFD approach, the mesh points for each discipline are embedded in a parametric volume. The control points that define the parametric volume are then manipulated to modify the embedded mesh points to obtain smooth changes to the discipline-level meshes. The disadvantage of the FFD approach is that the initial source geometry representation and the final geometry representation are not the same. However, the FFD approach is very flexible and can be applied to any mesh without knowledge of the underlying geometric representation. Furthermore, the FFD approach can be used to obtain efficient and accurate derivatives of the mesh points with respect to the geometric design variables. Obtaining these derivatives efficiently and accurately is crucial for multidisciplinary gradient computation, and is prohibitively expensive to achieve with CAD-based approaches.

In the following section, we outline a systematic way to manipulate the FFD control points to obtain geometry changes for an aircraft wing. In particular, we include changes to the local twist angle, span, chord, thickness-to-chord ratio, dihedral and sweep. In this work, we use three-dimensional B-spline volumes as the FFD volumes. However, the control point manipulation scheme presented here could be extended to other parametric volumes, such as radial basis function volumes. In the proposed scheme, geometric modifications are applied to the initial set of FFD control points, $\mathbf{p}_{ijk} \in \mathbb{R}^3$, to obtain the final set of control points, $\mathbf{P}_{ijk} \in \mathbb{R}^3$, where all coordinates are given in a global Cartesian reference frame. Chord, span and thickness-to-chord ratio are modified through an anisotropic scaling of the geometry along different directions, while twist, dihedral and sweep changes are applied in a consistent manner that avoids self-intersecting surfaces for large changes to sweep and dihedral and moderate changes to twist.

To apply these changes in a consistent manner, we employ a series of unit vectors that define a span-wise direction, \mathbf{t}_s , a chord-wise direction, \mathbf{t}_c , and a vertical direction, \mathbf{t}_v . In addition, we also employ a series of reference points, $\mathbf{r}_n \in \mathbb{R}^3$, for $n = 1, \dots, N$, connected by line segments. The geometry modification is divided into two steps: first, the geometry changes are applied to the reference points, second the location of the initial FFD points, \mathbf{p}_{ijk} , relative to the initial reference line segments is used to determine the final position of the FFD control points. The geometric variables are split into two groups: those given for each span-wise segment, and those given at each span-wise station. The geometric variables given for each segment consist of the scaling along the span-wise direction, are the s_n , dihedral, Γ_n , and sweep, Λ_n , while the geometric variables given for each span-wise station consist of the twist, θ_n , chord-wise scaling, c_n , and vertical scaling, v_n .

The following rotation matrix is used extensively in the proposed FFD manipulation scheme:

$$\mathbf{C}(\mathbf{a}, \varphi) = \cos \varphi \mathbf{I} + (1 - \cos \varphi) \mathbf{a} \mathbf{a}^T - \sin \varphi \mathbf{a}^\times,$$

where $\mathbf{a} \in \mathbb{R}^3$ is a unit vector such that $\mathbf{a}^T \mathbf{a} = 1$, and φ is the angle of rotation about the unit vector \mathbf{a} [30]. Note that this rotation matrix is defined such that the components of the transformed vector are expressed in the transformed reference frame.

In the proposed scheme, the geometric changes are first applied to the reference line segments. The difference between adjacent reference line points is denoted, $\mathbf{a}_n = \mathbf{r}_{n+1} - \mathbf{r}_n$. The reference line segment is modified in the following manner: first, the dihedral is applied, followed by a sweep modification and finally by a span scaling operation. These operations can be written as follows:

$$\mathbf{A}_n = s_n \mathbf{C}(\mathbf{b}, \Lambda_n)^T \mathbf{C}(\mathbf{t}_c, \Gamma_n)^T \mathbf{a}_n,$$

where $\mathbf{b} = \mathbf{C}(\mathbf{t}_c, \Gamma_n)^T \mathbf{t}_v$ is the vertical direction vector rotated through the dihedral angle. The final reference point locations, \mathbf{R}_n , are determined from applying the update

$$\mathbf{R}_{n+1} = \mathbf{R}_n + \mathbf{A}_n, \quad (29)$$

with $\mathbf{R}_1 = \mathbf{r}_1$, for $n = 1, \dots, N - 1$.

The twist axis, \mathbf{t}_θ , which defines the axis about which the twist rotation is applied, is determined by projecting the segment direction, \mathbf{A}_k , onto the plane defined by the span axis, \mathbf{t}_s , and vertical axis, \mathbf{t}_v , as follows:

$$\mathbf{t}_\theta = \frac{(\mathbf{t}_s \mathbf{t}_s^T + \mathbf{t}_v \mathbf{t}_v^T) \mathbf{A}_n}{\|(\mathbf{t}_s \mathbf{t}_s^T + \mathbf{t}_v \mathbf{t}_v^T) \mathbf{A}_n\|_2}. \quad (30)$$

To obtain the final geometry, the vertical axis and the chord axis are scaled and rotated based on the values of the twist, dihedral, chord and vertical scaling. In the final geometry, the modified vertical and chord axes are denoted \mathbf{v}_n and \mathbf{c}_n , respectively. These vectors are defined for each segment as follows:

$$\begin{aligned} \mathbf{c}_1 &= c_1 \mathbf{C}(\mathbf{t}_s, \theta_1)^T \mathbf{t}_c, & \mathbf{v}_1 &= v_1 c_1 \mathbf{C}(\mathbf{t}_s, \theta_1)^T \mathbf{t}_v, \\ \mathbf{c}_n &= c_n \mathbf{C}(\mathbf{t}_\theta, \theta_n)^T \mathbf{t}_c, & \mathbf{v}_n &= v_n c_n \mathbf{C}(\mathbf{t}_\theta, \theta_n)^T \mathbf{C}(\mathbf{t}_c, \tilde{\Gamma}_n) \mathbf{t}_v, \end{aligned}$$

where, $\tilde{\Gamma}_n = 1/2 (\Gamma_n + \Gamma_{n+1})$, $\tilde{\Gamma}_N = \Gamma_N$.

After the final reference line locations and the transformed chord and vertical axes, \mathbf{c}_n and \mathbf{v}_n , have been calculated, the final FFD control point locations are determined based on the values of the following projections:

$$\begin{aligned} u_s &= \frac{\mathbf{t}_s^T}{\mathbf{t}_s^T \mathbf{a}_n} (\mathbf{p}_{ijk} - \mathbf{r}_n), \\ u_c &= \mathbf{t}_c^T (\mathbf{p}_{ijk} - \mathbf{r}_n - u_s \mathbf{a}_k), \\ u_v &= \mathbf{t}_v^T (\mathbf{p}_{ijk} - \mathbf{r}_n - u_s \mathbf{a}_k), \end{aligned}$$

where u_s is the projection onto the span direction, u_c is the projection onto the chord direction and u_v is the projection onto the vertical direction. If $0 \leq u_s < 1$, then the following update is applied:

$$\mathbf{P}_{ijk} = \mathbf{R}_n + u_s \mathbf{A}_n + u_c ((1 - u_s) \mathbf{c}_n + u_s \mathbf{c}_{n+1}) + u_v ((1 - u_s) \mathbf{v}_n + u_s \mathbf{v}_{n+1}). \quad (31)$$

If $u_s < 0$ or $u_s \geq 1$, then \mathbf{P}_{ijk} is unmodified by the segment. Figure 5 shows the FFD volume points and reference points and line segments for an initial straight wing, and a modification of geometry to a swept C-wing with taper and a crank.

IV. Aerostructural design problem formulation

In this work we consider the design of both a metallic and a composite wing. We compare the results of a series of optimizations to understand the trade-offs that occur between the aerodynamic drag and structural weight and how the use of composite structures shifts where these trade-offs occur.

The geometry of all cases presented in the following sections is based on the Boeing 777-200ER aircraft. The main focus of the design problems is the wing geometry and structure. However other geometric data and aircraft

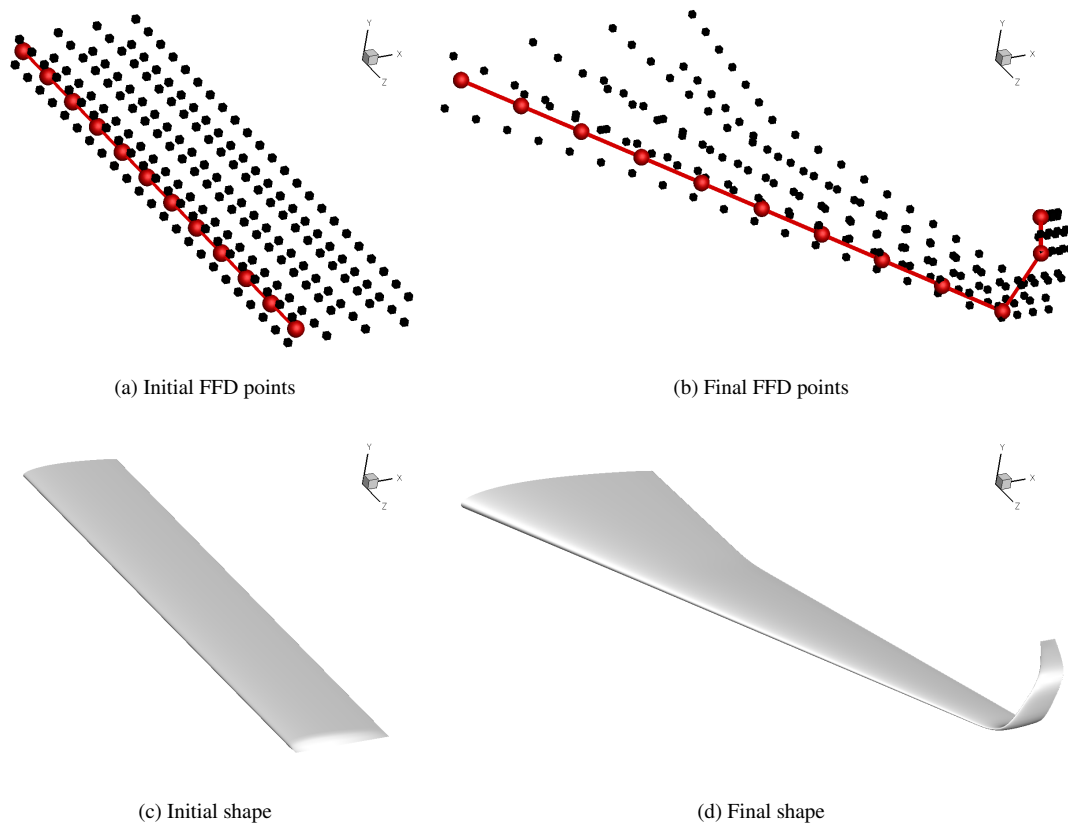


Figure 5: A geometry modification from an initial straight wing to a swept C-wing with taper and a crank.

parameters are required to obtain reasonable drag estimates and fuel consumption predictions. These additional data are obtained in part from the Boeing 777-200ER publicly available performance data. All additional data required for the design problem are listed in Table 1. The initial wing span is 60.9 m with a taper ratio of 0.2, a root chord of 13.2 m at the symmetry plane, and a sweep of 30°. The wing crank occurs at 30% of the semi-span. The initial wing is a linear loft of RAE2822 airfoil sections without twist or dihedral. The wing structure consists of 44 ribs, and three spars: a leading edge spar, a trailing edge spar, and landing gear sub-spar. The ribs before the wing crank are placed chord-wise, while the ribs in the outer-portion of the wing are placed perpendicular to the rear spar. The landing gear sub-spar extends from the nominal landing gear join to the symmetry plane.

Table 2 lists the material data for both the metallic, aluminum wing and the composite wing. In both cases we apply a factor of safety of 0.8 to the material data in all calculations. The aluminum material properties are representative of typical Al-7075 properties, while the composite properties are typical of a high-strength carbon-epoxy system.

Parameter	Value	Units
Cruise Mach number	0.84	
Operating empty weight (OEW)	138 100	kg
Design payload	40 040	kg
Initial wing offset mass	30 000	kg
Fixed design mass (m_{fixed})	148 140	kg
Thrust-specific fuel consumption (TSFC)	0.53	lb/(lb hr)
Geometric parameter	Value	Units
Reference area (S_{ref})	423	m ²
Initial span (b)	60.9	m
Horizontal stabilizer area (S_{hstab})	90	m ²
Horizontal stabilizer MAC (c_{hstab})	4	m
Vertical stabilizer area (S_{vstab})	90	m ²
Vertical stabilizer MAC (c_{vstab})	5	m
Nacelle area (S_{nacelle})	150	m ²
Nacelle reference length (c_{nacelle})	6	m
Fuselage area (S_{fus})	1036	m ²
Fuselage reference length (S_{fus})	63	m

Table 1: A summary of additional data used during the analysis of the aircraft.

Parameter	Value	Units	Parameter	Value	Units
Aluminum material data					
E	70.0	GPa	ν	0.3	
σ_{YS}	380	MPa	ρ	2780	kg/m ³
Composite material data					
E_1	128	GPa	E_2	11	GPa
G_{12}	4.5	GPa	G_{13}	4.5	GPa
G_{23}	3.2	GPa	ν_{12}	0.25	
X_t	1170	MPa	X_c	1120	MPa
Y_t	40	MPa	Y_c	170	MPa
S	48	MPa	ρ	1522	kg/m ³

Table 2: The metallic and composite material properties

A. Design objective

In order to assess the effect of different design objectives, we use a weighted combination of the takeoff gross-weight (TOGW) and fuel burn (FB) over a 8000 nm mission. The rationale for this choice is that the fuel burn, at fixed cruise

Mach number, is a significant contributor to the direct operating costs of the aircraft, while takeoff gross-weight is a rough surrogate for the overall aircraft acquisition cost. The TOGW objective places larger emphasis on minimizing the structural weight of the aircraft, while the fuel burn objective places more emphasis on increasing the lift-to-drag ratio. By examining a weighted objective we are able to assess a series of feasible designs that place an increasing importance on the aerodynamic performance of the aircraft.

The mission profile used in the analysis is illustrated in Figure 6. Here the mission is split into three equal segments at different altitudes with a constant cruise Mach number of $M = 0.84$. The total mission fuel weight is estimated based on an analysis of each of these cruise segments. Fuel consumption during taxi, takeoff, climb and landing is not included in the analysis. Over this long range mission, the fuel consumption over these secondary segments is a small fraction of the overall fuel consumption, and a more detailed analysis of the off-design performance would be required to accurately determine these contributions [31].




Figure 6: The mission profile illustrating the different cruise segments.

The mass ratio of each segment is calculated based on the Breguet range equation

$$\frac{W_i}{W_{i+1}} = \exp\left(\frac{R_i c_i}{V_i (L/D)_i}\right), \quad (32)$$

where c_i , R_i and V_i are the thrust specific fuel consumption (TSFC), range, and cruise speed for the i^{th} cruise segment. The lift-to-drag ratio $(L/D)_i$ is evaluated at the mid-point of the cruise segment at the aircraft weight $W_{i+1/2}$, as illustrated in Figure 6. The lift-to-drag ratio varies over the different cruise segments due to aerostructural effects produced by changes in aircraft weight and the inertial fuel load distribution. Altitude changes also change the lift-to-drag ratio through Reynolds number effects. The atmospheric conditions at each cruise segment are calculated using the ICAO extended standard atmosphere. In order to simplify the analysis, we maintain a constant TSFC over all altitudes, $c_i = 0.53 \text{ lb}/(\text{lb hr})$.

The fuel consumption over the entire mission is given by the difference between the takeoff gross-weight and the landing weight of the aircraft:

$$\text{FB} = W_0 - W_N = W_N \left(\frac{W_0}{W_N} - 1 \right), \quad (33)$$

where W_0/W_N can be obtained using Equation (32) as follows:

$$\frac{W_0}{W_N} = \prod_{i=0}^{N-1} \frac{W_i}{W_{i+1}}$$

The objective we investigate is a weighted average of the fuel burned during the maximum range mission, and the TOGW of the maximum range mission. This objective can be written as follows:

$$f(\mathbf{x}) = \beta \text{FB} + (1 - \beta) \text{TOGW} = W_N \left(\frac{W_0}{W_N} - \beta \right) \quad (34)$$

where β is a weighting parameter such that at $\beta = 0$, results in a takeoff gross-weight objective, while $\beta = 1$ results in a mission fuel burn objective. At intermediate values of $\beta \in (0, 1)$, the objective is to minimize a weighted combination of the fuel burn and takeoff gross-weight. The effect of the β parameter is investigated in Section V.

B. Structural design conditions

In addition to the mission analysis presented above, we also perform aerostructural analyses at several off-design conditions in order to enforce structural constraints in compliance with parts of the FAR Part 25 regulations. In particular, we impose constraints at the indicated points on the V-n diagram shown in Figure 7 at 100% of the mission fuel load and at 10% of the mission fuel load, both at an altitude of 20 000 ft. The different fuel conditions result in different inertial relief which in turn results in different wing-load distributions. The Mach number at the 2.5g and -1g dive conditions is fixed at $M = 0.88$ with an angle of attack variable added to ensure that the lift is 2.5 times the total mass of the aircraft. For the 2.5g stall condition it is necessary to add both an angle of attack variable, to ensure sufficient lift and an air speed variable to ensure that $C_L = C_{Lmax}$ at the flight condition. The C_{Lmax} of the clean wing is determined by using the Valarezo maximum lift condition presented in Section III.

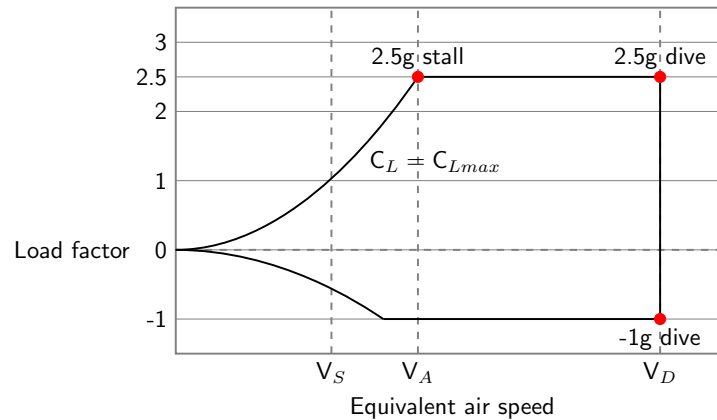


Figure 7: The V-n diagram with the diagrams for a given fuel mass.

C. The design parametrization

The design variables can be partitioned into three groups: aerodynamic design variables, structural design variables and geometric design variables. It should be emphasized that in the context of aerostructural analysis and design optimization, the aerodynamic design variables affect the structural analysis and the structural design variables affect the aerodynamic performance. Here we consider two design problems: the design of a metallic wing and the design of a composite wing under identical sets of operating conditions and nearly identical sets of design constraints. The geometric and aerodynamic design variables are common to both the metallic and composite design problems, while the structural design parametrization is significantly different.

The aerodynamic design variables consist of the angle of attack at each of the 9 flight conditions and the air speed at the two C_{Lmax} conditions. In addition, we use 4 fuel mass variables, 3 that represent the mid-points of the flight segments illustrated in Figure 6, and an additional fuel burn variable that represents the full fuel load.

The geometric design variable parametrization is illustrated in Figure 8. For the geometric parametrization, we use 10 reference point locations positioned from the root to the tip at the trailing edge of the wing. The first 3 reference points are positioned uniformly from the wing root to the wing crank, while the remaining points are positioned uniformly from the wing crank to the wing tip. The chord scaling variables are linked such that $c_n = c_1$, for $n = 2, \dots, 10$. The span scaling variables are also linked such that $s_n = s_1$ for $n = 2, \dots, 9$. We also use the vertical scaling variables over the range, $0.75 \leq v_n \leq 1.25$. Since the initial airfoil section has a t/c ratio of approximately 12%, these bounds ensure that the t/c ratio varies between 9% and 15%. A series of linear constraints are imposed on the vertical scaling variables such that the variables v_1, v_3, v_{10} , are independent, while all remaining vertical scaling variables are interpolated linearly between these values. Finally, we use 9 twist design variables, θ_n , with the root-twist fixed, $\theta_1 = 0$.

The structural design parametrization for the metallic wing consists of the thickness and panel geometry variables shown in Figure 4. Each independent panel has 3 thicknesses and 3 geometric variables, where the stiffener pitch, b , is fixed for all panels on the upper and lower surfaces, respectively. To reduce the number of variables in the design problem, the designs of adjacent panels are linked in groups of two. As a result there are 138 thickness variables and 138 panel geometric variables. In addition to the panel variables, the thicknesses of spars, ribs and the leading edge for each segment are also included as design variables, resulting in an additional 186 thickness variables. There are also

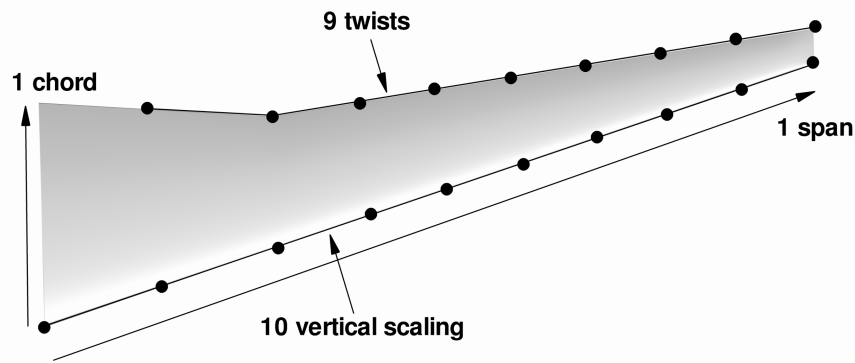


Figure 8: A summary of the geometric design variables

114 fuel mass variables for the inertial fuel loading conditions and 46 panel-length design variables required for the buckling calculations. As a result, there are 462 thickness and panel geometry variables, and 160 consistency design variables, which include the fuel inertial mass and panel-length variables, resulting in 622 structural design variables.

For the composite wing design parametrization, the laminate parametrization variables described in Section III are included in the design problem, in addition to the thickness and panel geometry design variables described above that are also used for the metallic wing. The number of variables in the laminate parametrization depends on the initial thickness distribution selection. If a total of n layers are required then $4n$ design variables are added to the design problem. Typically there are between 200 and 500 plies in the initial thickness distribution. As a result, an additional 800 to 2000 design variables are required for the laminate parametrization. In total there are $658 + 4n$ design variables, where n are the number of plies included in the laminate parametrization, with $n = 0$ for the metallic wing.

D. Constraint formulation

Nonlinear constraints		Linear constraints	
Fuel burn compatibility	4	t/c linearity	7
Lift constraints	9	Thickness variation	126
2.5g dive KS	10	Stiffener height variation	44
-1g dive KS	10	Spar variation	86
2.5g stall KS	10	Stiffener dimension	184
Valarezo C_{Lmax}	2	Ply-identity summation	n
Panel geometry compatibility	46	Balanced laminate condition	m
Fuel volume	1	10% ply content	p
Fuel mass-per-area compatibility	114	Ply contiguity constraint	q
Total	206		$477 + n + m + p + q$

Table 3: Summary of the constraints for the aerostructural problem

The constraints for the aerostructural design problem are summarized in Table 3. There are at least 683 constraints: 206 nonlinear constraints and at least 477 linear constraints, depending on the number of plies in the parametrization. There are four fuel burn compatibility constraints corresponding to the fuel burn calculations for the mid-points of the flight segments illustrated in Figure 6, and the total mission fuel. There are a total of 30 KS failure and buckling constraints at 6 separate flight conditions. At each flight condition there are 3 KS failure constraints: one aggregated over each of the top skin, bottom skin, and spars and ribs, and 2 KS buckling constraints: one aggregated over each of the top and bottom skins. For all cases, we use an aggregation parameter of $\rho = 50$. The Valarezo criterion described above is also applied for the two 2.5g stall conditions. The panel geometry compatibility constraints ensure that the physical panel lengths correspond to the variables that represent the dimensions of the panels. The fuel

volume constraint ensures that the total mission fuel can fit inside the spar box and the full mass-per-area compatibility constraints ensure that the correct inertial fuel load is applied to the structure.

The linear constraints consist of the constraints to impose the piecewise linearity of the t/c distribution. The thickness variation, stiffener height variation and spar variation constraints ensure that the change in thickness and spar height do not exceed 5 mm, or 1 cm between adjacent panels, respectively. Finally, a series of linear constraints are imposed on the spar height and stiffener width to ensure that they remain within reasonable bounds. The size of the remaining linear constraints that are applied to the laminate parametrization variables depends on the initial ply distribution.

E. Summary of the proposed studies

The aerostructural optimization studies presented in the following sections can be written in the following manner:

$$\begin{aligned}
 & \text{minimize} && \beta \text{FB} + (1 - \beta) \text{TOGW} \\
 & \text{w.r.t.} && \mathbf{x} \\
 & \text{governed by} && \mathbf{R}(\mathbf{q}_j, \mathbf{x}) = 0 && j = 1, \dots, 9 \\
 & \text{s.t.} && L_j(\mathbf{q}_j, \mathbf{x}) = n_j g(m(\mathbf{x}) + m_{\text{fixed}}) && j = 1, \dots, 9 \\
 & && \text{KS}_{\min}(\Delta C_{p_{\text{margin}}}(\eta_k), \rho) = 0 && j = 7, 8 \\
 & && \text{KS}_k(F_{\text{KS}}(\boldsymbol{\sigma}), 30), 50) \leq 1 && j = 4, \dots, 9 \quad k = 1, \dots, 3 \\
 & && \text{KS}_k(B(N_x, N_{xy}), 50) \leq 1 && j = 4, \dots, 9 \quad k = 1, 2 \\
 & && h(\mathbf{x}) \geq 0
 \end{aligned}$$

where \mathbf{x} are the design variables, β is the weighting parameter, n_j is the load factor for each of the 9 flight conditions, and $q_j = 1/2\rho_j V_j^2$ is the dynamic pressure. Here, j indexes the flight condition, where $j = 1, \dots, 3$ correspond to the mission flight segments, $n_1 = n_2 = n_3 = 1$, $j = 4, 5$ correspond to the 2.5g maneuver condition with full fuel load and 10% fuel load, respectively, $n_4 = n_5 = 2.5$, $j = 6, 7$ corresponds to the -1g maneuver condition with full fuel load and 10% fuel load, respectively, $n_6 = n_7 = -1$, and $j = 7, 8$ corresponds to the 2.5g stall condition with full fuel load and 10% fuel load, respectively, $n_8 = n_9 = 2.5$. Note that the KS constraints for material failure and buckling are applied only at the 2.5g and -1g maneuver conditions. Finally, $h(\mathbf{x})$ represents the remainder of the constraints listed in Table 3.

We use a finite-element structural model with 46 360, 3rd order MITC9 shell elements with just over 1.088 million degrees of freedom. There are a total of 46 local panel buckling models, each requiring two buckling calculations: the axial buckling load and the shear buckling load. The buckling calculation for both cases require approximately 11 seconds of CPU time each, but are distributed across all processors assigned to the aerostructural optimization problem. The aerodynamic model consists of 4200 surface panels and 60-streamwise wake panels.

All aerostructural optimization cases are distributed across 72 processors, which are subdivided into 3 aerostructural analysis and gradient evaluation groups of 24 processors each. The 9 flight conditions are assigned to the 3 groups: 1 cruise flight condition, and 2 off-design structural conditions per aerostructural group. This division divides the computational effort for the objective and constraint evaluation and gradient evaluation into roughly equal portions. The solution of all the aerostructural flight conditions and the evaluations of the objective and constraints requires roughly 3 minutes and 30 seconds of wall time, while the gradient evaluation of all constraints and the objective requires 5 minutes and 15 seconds of wall time.

V. Results

A. Structural optimization results

Before beginning an aerostructural optimization, we solve a preliminary structural mass-minimization problem with fixed aerodynamic loads and use the structural design as the starting point for the aerostructural problem. The structural design obtained from this preliminary structural optimization provides a much better starting point than a constant-thickness wing or other arbitrary design. Aerostructural optimizations started from the structurally-optimized design typically require fewer optimization iterations than an initial arbitrary structural design point. Furthermore, the structurally optimized wing satisfies the structural constraints, which makes for a useful comparison between the initial and final aerostructural results.

To generate the aerodynamic forces for the structural mass-minimization problems, we use a metallic wing with a constant thickness distribution, and obtain a converged aerostructural solution, for each of the flight conditions

described above using the approximate Newton–Krylov method. We use the aerodynamic forces from this converged aerostructural solution to obtain the fixed structural loads. In the aerostructural case, changes in the stiffness of the wing would produce changes in the deflection and, in turn, different aerodynamic loads. In the context of this optimization, these effects are ignored. For both the metallic and the composite wing, the constraint formulations follow the formulation presented in Table 3, without the fuel burn, lift, Valarezo or fuel volume constraints.

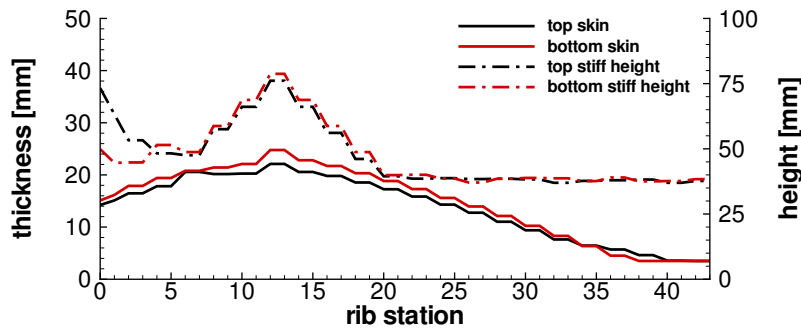

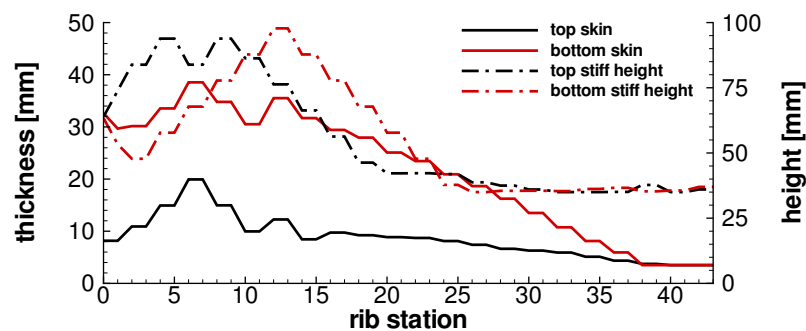


Figure 9: The thickness distribution for the metallic wing

We solve the metallic wing structural optimization problem on 48 processors on the General Purpose Cluster (GPC) at SciNet [32]. Each node of the GPC is an Intel Xeon E5540 with a clock speed of 2.53GHz, with 16GB of dedicated RAM and 8 processor cores. All calculations are performed on the SciNet system with the same configuration. The total mass of the structurally-optimized metallic wing is 32 660 kg. The thickness distribution for the upper and lower wing skins and the stiffener heights are shown in Figure 9. The optimization requires 256 objective and constraint evaluations and 159 gradient evaluations. The wall time to update the panel-level models including all buckling calculations is approximately 22 seconds, while the wall time to set up and solve the global finite-element model is 14 seconds. The wall time to evaluate the total derivatives of the objective and constraints is approximately 15 seconds. As a result, the total structural optimization time is roughly 3 hours 15 minutes.



(a) Ply distribution 1



(b) Ply distribution 2

Figure 10: The thickness distribution for the composite wing

For the composite wing optimization case, we use the optimized metallic wing to obtain an initial estimate of the number of plies in each component. We label this distribution of plies, “Ply distribution 1”. In order to adjust the thickness along the wing span, we use a ply-blending scheme in which the initial thickness distribution is obtained by adding or removing the outer-most plies from the laminate. As a result, the middle ply in the laminate runs from the wing root to the wing tip, while the outer-most plies extend only over the thickest portion of the wing. As described in Section III, we solve a series of optimization problems for an increasing sequence of penalty parameter values with $\gamma_1 = 0$, and $\gamma_p = 10^{-5}2^{p-2}$, for $p > 1$, where p indicates the optimization problem sequence. In each case, we start the next optimization problem from the previous solution. This sequence of problems forces the infeasibility with respect to the spherical constraints (18) to zero, and produces a lamination sequence in which only one ply-identity variable is active in each layer [23].




Figure 11: Lamination sequences for the structural-only optimization problem. Only the top half of the symmetric laminate is shown. The stacking sequences are split where T2 is stacked on top of T1 and B3 is stacked on top of B2, which is stacked on top of B1.

The mass of the wing obtained from the optimization problem outlined above using Ply distribution 1 is 21 298 kg. The optimized thickness distribution is shown in Figure 10a. This thickness distribution is significantly different than the metallic distribution of plies, shown in Figure 9. As a result, there is a significant discrepancy between the number of plies in the design and the physical number of plies required. In order to assess the severity of this discrepancy, we obtain the number of plies using the optimized thickness distribution shown in Figure 10a. We label this new

distribution of plies “Ply distribution 2” and perform the same optimization with this new ply distribution. Figure 10b shows the results of the structural optimization using Ply distribution 2. The new optimized mass of the wing is 21 177 kg. There is only a 121 kg difference between the wing masses, despite the original ply miss-match. Note also that discrepancy between the thickness distributions for Ply distribution 2, Figure 10b, and Ply distribution 1, Figure 10a, is much smaller than between the metallic distribution and the Ply distribution 1 results. This suggests that by updating the ply thickness distribution, and performing a second optimization, we can come much closer to the actual ply thickness distribution. However, these results also suggest that the weight will not vary significantly between the two designs. Based on this assessment, we do not perform a secondary optimization step for the aerostructural optimization results.

Figure 11 shows the stacking sequence results for the structural mass-minimization problem for both ply distributions. The top and bottom lamination sequences are too large to fit on a single page, and have been split, where the sequences labeled T2 is stacked on top of the sequence labeled T1, and the stacking sequence B3 is stacked on top of B2 which is stacked on top of B1. The balanced ply constraint, 10% ply content constraint, and 4-ply adjacency constraint can be checked by inspection. The Ply distribution 2 case contains 4 plies which have not fully converged. These plies have intermediate $\pm 45^\circ$ ply-identity variables while all other variables in these layers are zero. The two designs share some common attributes. Specifically, the top and bottom stringer are predominantly 0° plies, with the minimum ply content. The top skin has more $\pm 45^\circ$ plies than the bottom skin, which has more 0° plies. These results suggest that the ply distribution in the different components will be similar between designs with different initial ply distributions, even if the number of plies is incorrect.

B. Metallic and composite wing aerostructural optimization

In this section we present a series of aerostructural optimization results obtained with $\beta = 1, 0.875, 0.75, 0.5$ and 0 , for the metallic and composite wings. This uneven distribution of weights was selected based on preliminary optimization studies which found that the designs vary most rapidly over the range $\beta \in [0.5, 1]$. This should be expected, given that the fuel burn is roughly one third of the takeoff gross-weight of the aircraft.

For all optimization studies we use the nonlinear optimization code SNOPT [25], through the Python interface of the pyOpt optimization package [26]. For all cases, we use an optimality tolerance of 10^{-5} and a feasibility tolerance of 10^{-5} . The aerostructural optimizations take between 32 hours and 48 hours on 72 processors.

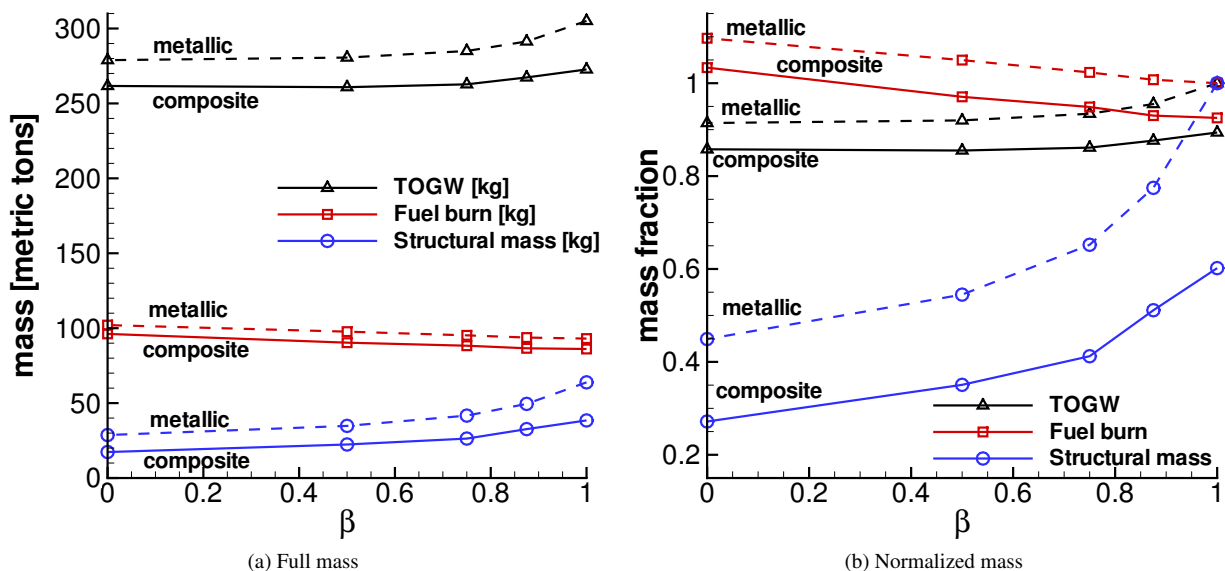


Figure 12: The TOGW, fuel burn and structural wing mass as a function of the parameter β for the metallic and composite wings. Note that the mass fraction results are normalized to the fuel burn result for the metallic wing.

Figure 12 shows the optimized takeoff gross-weight, fuel burn and structural wing weight of the metallic and composite aircraft as a function of the parameter β . For the metallic wing, the TOGW and fuel burn vary between 278 888 kg and 102 062 kg, respectively for $\beta = 0$, and 305 056 kg and 93 055 kg, respectively for $\beta = 1$. For the composite wing, the TOGW and fuel burn vary between 261 664 kg and 96 176 kg, respectively for $\beta = 0$, and

272 679 kg and 86 101 kg, respectively for $\beta = 1$. The fuel burn designs place a much higher importance on the aerodynamic performance with a structural weight and TOGW penalty. The structural weight for the metallic wing varies from 28 686 kg for $\beta = 0$ to 63 861 kg for $\beta = 1$, while the structural weight for the composite wing varies from 17 348 kg for $\beta = 0$ to 38 438 kg for $\beta = 1$.

Figure 12b shows the variation of the TOGW, fuel burn, and structural wing weight normalized to the metallic fuel burn minimization result. Clearly the TOGW, fuel burn and structural weight of the composite aircraft are less than the equivalent metallic design for all values of β . The weight of the composite wings are between 34% and 40% lighter than the equivalent metallic wings. Despite this large structural weight savings, the difference in fuel burn between any pair of metallic and composite designs is between 5% and 8%. The TOGW of the composite aircraft are between 6% and 11% lower than the TOGW of the equivalent metallic aircraft.




Figure 13: The span and average L/D ratio for the metallic and composite wings as a function of the parameter β .




Figure 14: The wing wetted area and aspect ratio for the metallic and composite wings as a function of the parameter β .

As the parameter β is increased from 0, the structural weight and TOGW of both the metallic and composite aircraft increase significantly. The benefit of this large increase in wing weight is significantly higher aerodynamic performance. Figure 13, Figure 14 and Figure 15 show the change in aerodynamic performance and geometric variables as a function of β for both the metallic and composite cases. Figure 13 shows the average L/D ratio from the three cruise flight segments shown in Figure 6, as well as the wing span as a function of the parameter β for both the metallic and composite designs. For the metallic wing, the average L/D improves from a value of 19.3 at $\beta = 0$ to a value of 23.8 at $\beta = 1$. This 23% improvement is primarily achieved through an extension to the wing span, also shown in Figure 13a. For the composite wing, the average L/D improves by 18% from 19.2 at $\beta = 0$ to a value of




Figure 15: The span efficiency e during the first flight segment as a function of the parameter β for the metallic and composite wings

22.7 at $\beta = 1$. This improvement is also achieved through a span extension. Despite the lower average L/D ratio, the composite aircraft burns less fuel than the metallic aircraft. This poorer L/D performance is due to the lower C_L values required by the lighter composite aircraft. Due to the fixed flight segment altitudes, the aircraft cannot fly at their optimal L/D values. Permitting a variable altitude would result in a higher L/D for the composite aircraft. For comparison, Figure 15 shows the span efficiency of the metallic and composite designs. At the design for $\beta = 0.75$, the span efficiencies of the wings are nearly equal, while at all other points, the composite wing exhibits a better span efficiency than the metallic wings.

The increase in span, L/D and wing weight is also accompanied by an increase in the wetted area and the aspect ratio of the metallic and composite wings. Figure 14 shows the wetted area and aspect ratio as a function of the parameter β . For the metallic wing, the aspect ratio increases from 7.5 to 12.8, and the wetted wing area increases from 901 m² to 1034 m². For the composite wing, the aspect ratio increases from 8.3 to 13.7 while the wetted area increases from 829 m² to 1013 m². As before, the larger aspect ratios of the composite designs, and the larger wetted area of the metallic designs would suggest that the composite wings should exhibit superior aerodynamic performance. As discussed above, this is due to the sub-optimal cruise C_L values required by the lighter composite aircraft, which produces poorer L/D values than the equivalent metallic designs.




Figure 16: The planform of each of the metallic wings for $\beta = 1, 0.875, 0.75, 0.5,$ and 0 .

Figure 16 shows a comparison between the different wing planforms for the metallic wing with $\beta = 1, 0.875, 0.75, 0.5,$ and 0 . Figure 17 shows a comparison between the different wing planforms for the composite wing with $\beta = 1, 0.875, 0.75, 0.5,$ and 0 . The large changes in wing span, aspect ratio and wetted area are evident. The trend




Figure 17: The planform of each of the composite wings for $\beta = 1, 0.875, 0.75, 0.5,$ and 0 .

between the metallic and the composite wings are similar, where the largest discrepancy occurs for the wing spans of the $\beta = 0.875$ pair of cases, where the composite wing being 6% longer span than the metallic wing.

Figure 18 and Figure 19 show the aerostructural and jig twist distributions for the metallic designs at the different operating points as well as the sectional C_l distributions at all operating points for $\beta = 1, 0.875$ and $0.75,$ and 0.5 and $0,$ respectively. Figure 20 and Figure 21 show the aerostructural and jig twist distributions for the composite designs at the different operating points as well as the sectional C_l distributions at all operating points for $\beta = 1, 0.875$ and $0.75,$ and 0.5 and $0,$ respectively. In all cases there is a significant difference between the jig twist and the twist distributions at the cruise and maneuver conditions. Furthermore, for all cases the maneuver conditions exhibit more wash-out resulting in lower tip loads than the cruise conditions.

Figure 22 and Figure 23 show the top and bottom skin thickness distributions as well as the top and bottom stiffener height distributions for the metallic and composite wing design problems, respectively. All designs exhibit increasing thicknesses near the wing crank. This is due to the large running-loads carried in the top and bottom skins in this location. In all cases, as the parameter β increases, the structural thicknesses, and structural weight increase. This is due in part, to the larger spans of these designs. These larger spans require thicker structures and larger stiffeners to support the aerodynamic loads applied to these wings.

VI. Conclusions

In this paper, we presented a framework for aerostructural optimization of metallic and composite aircraft. In order to perform a fair comparison between the metallic and composite designs, we optimized the aircraft using an equivalent problem formulation where the primary difference in the formulations were the material properties and the detailed laminate parametrization variables. The extra design freedom provided by the detailed laminate parametrization and the superior strength-to-weight ratio of the composite enabled the composite wings to achieve spans up to 6% larger than the metallic wing spans. In addition, we found that the composite wing designs were between 34% and 40% lighter than the equivalent metallic wings. Due to this large structural weight savings, the composite aircraft designs exhibit a fuel burn savings of between 5% and 8% and a take-off gross-weight savings of between 6% and 11%.

Acknowledgement

The computations for the wing-box study presented within this paper were performed on the GPC supercomputer at the SciNet HPC Consortium at the University of Toronto. SciNet is funded by the Canada Foundation for Innovation, under the auspices of Compute Canada, the Government of Ontario, and the University of Toronto.

References

- [1] Jones, R. T., "The spanwise distribution of lift for minimum induced drag of wings having a given lift and a given bending moment," Tech. Rep. TN-2249, NASA, December 1950.

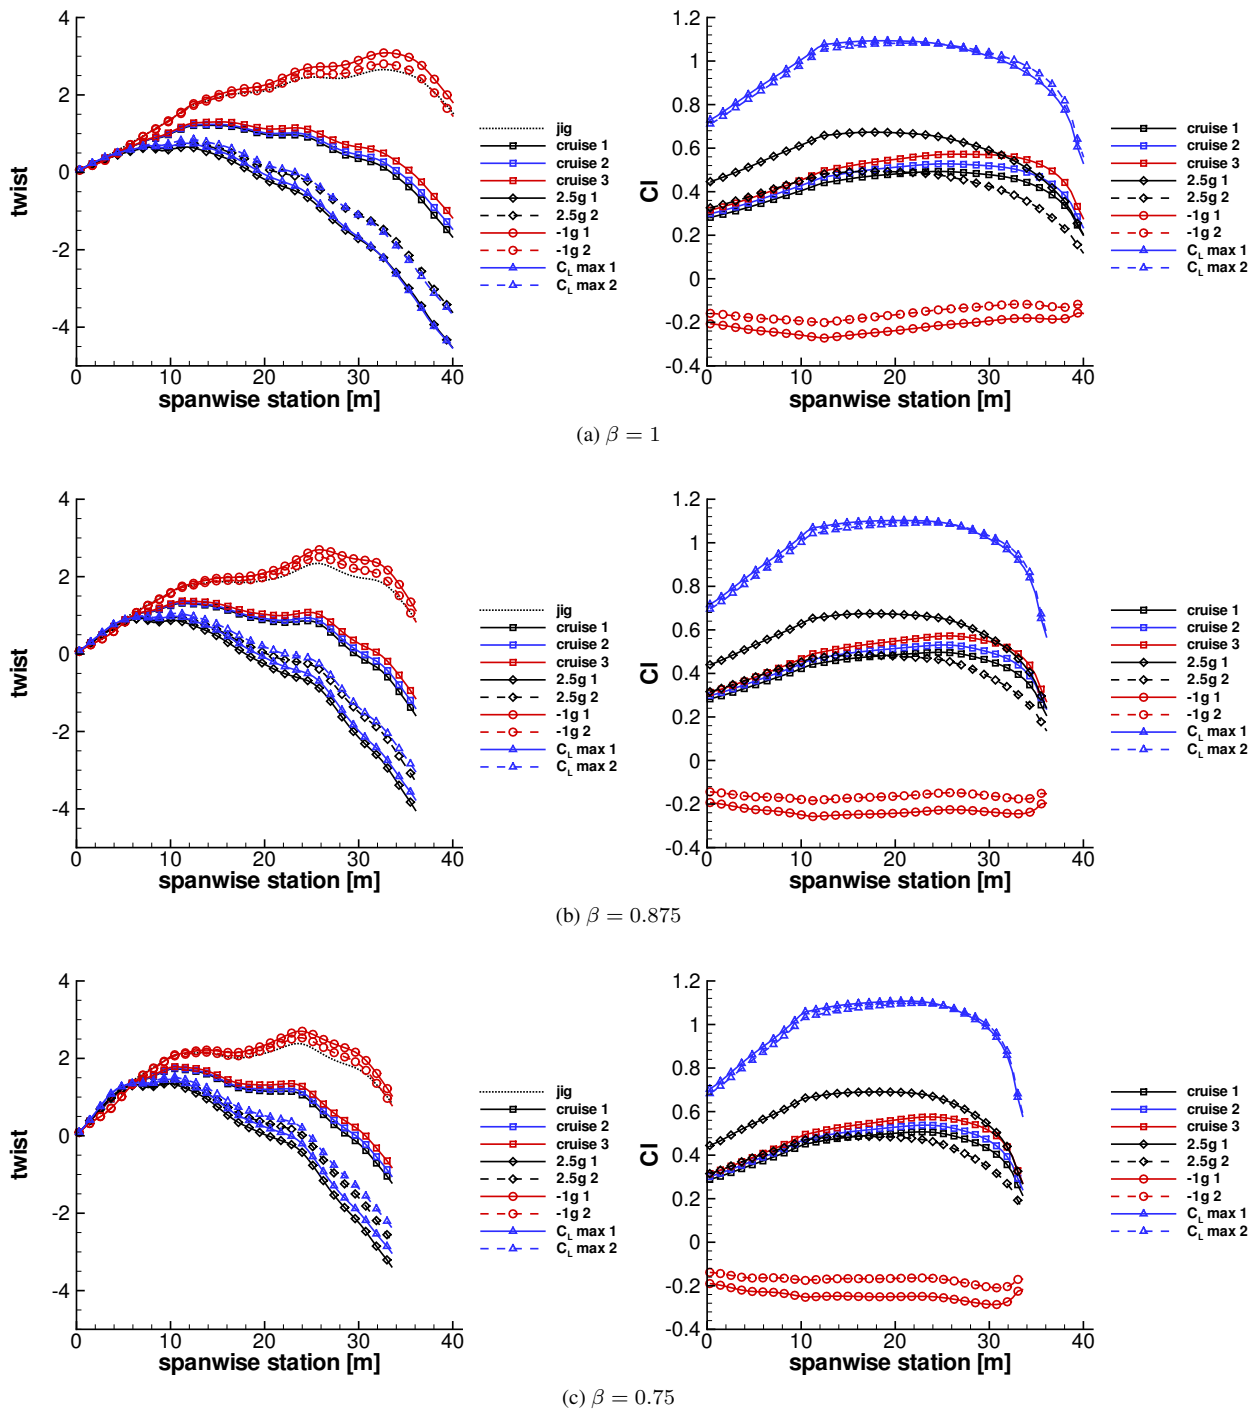


Figure 18: The twist and C_l distributions for the aerostructural optimization of a metallic wing with $\beta = 1, 0.875,$ and 0.75 .




Figure 19: The twist and C_l distributions for the aerostructural optimization of a metallic wing with $\beta = 0.5$, and 0.

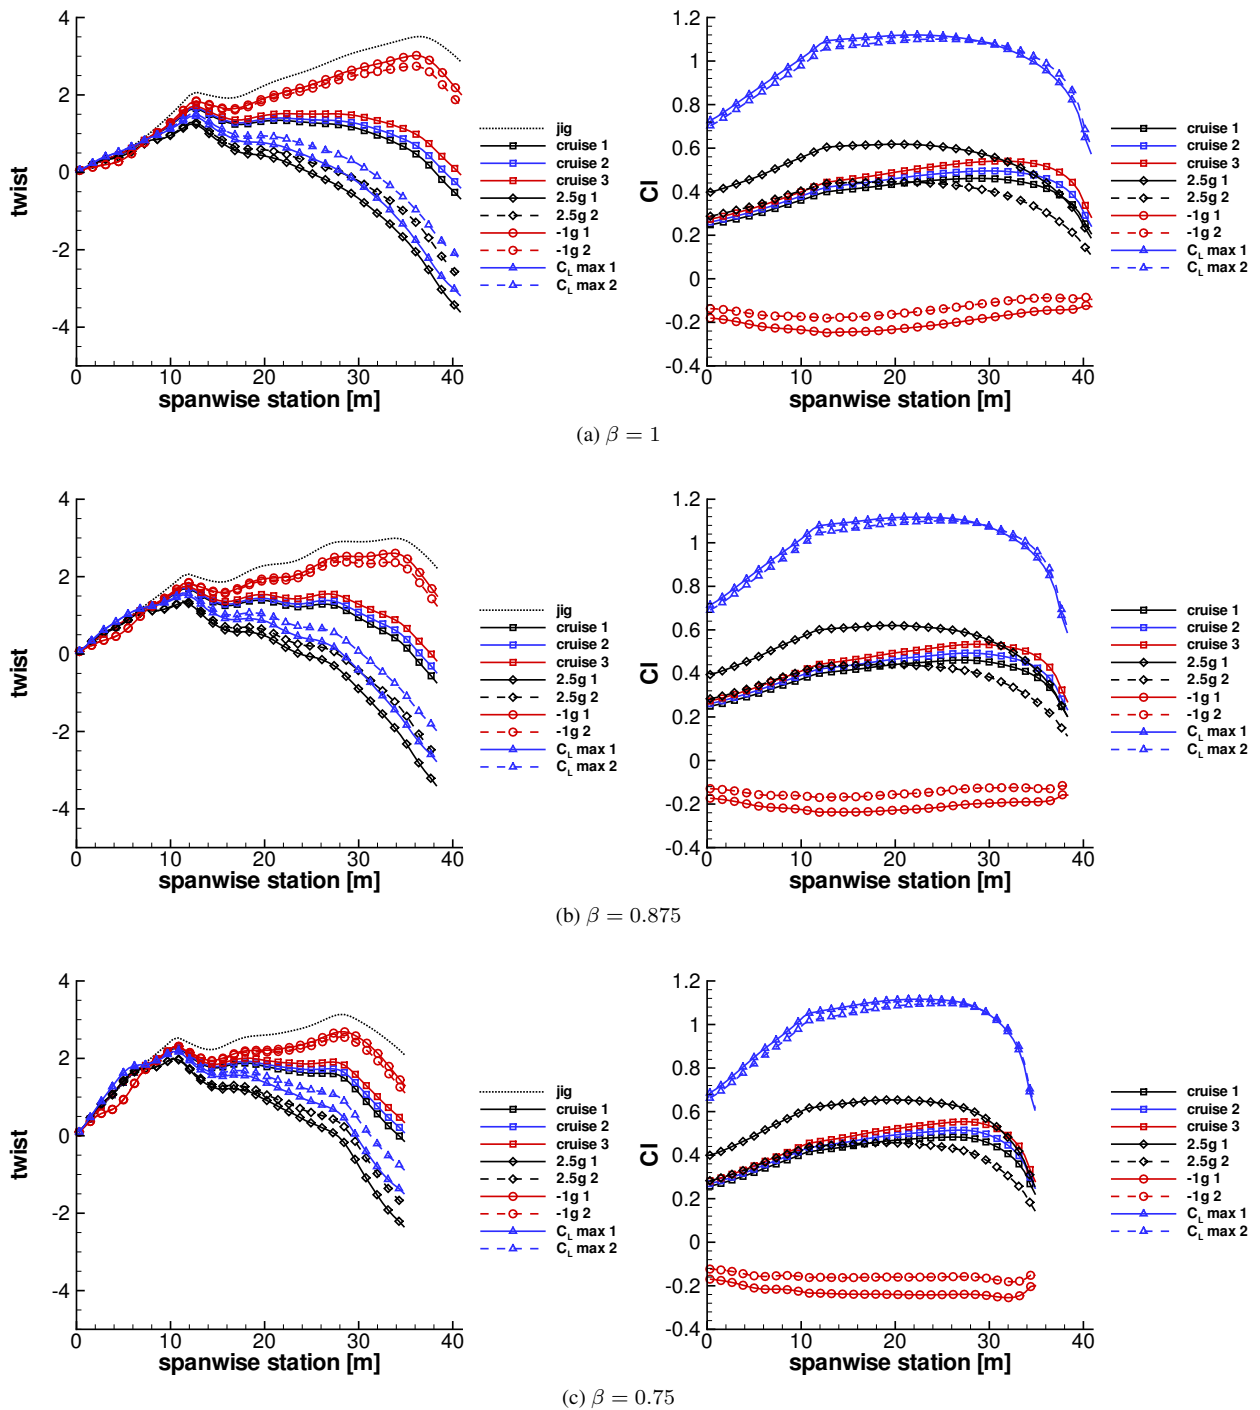


Figure 20: The twist and C_l distributions for the aerostructural optimization of a composite wing with $\beta = 1, 0.875,$ and 0.75 .




Figure 21: The twist and C_l distributions for the aerostructural optimization of a composite wing with $\beta = 0.5$, and 0.

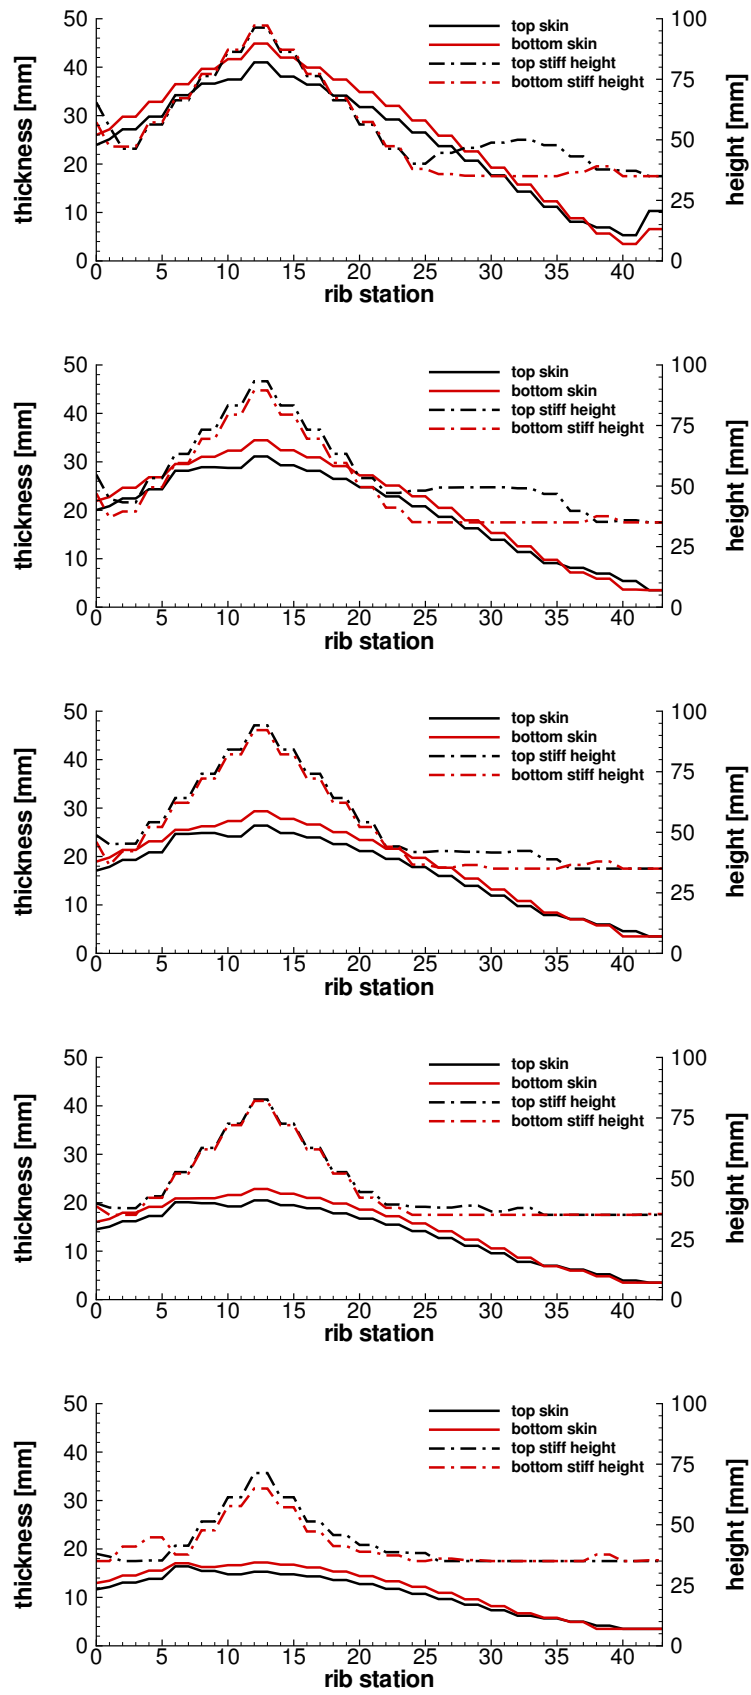


Figure 22: The skin thickness and stiffener height for the metallic wings for $\beta = 1, 0.875, 0.75, 0.5,$ and 0 .

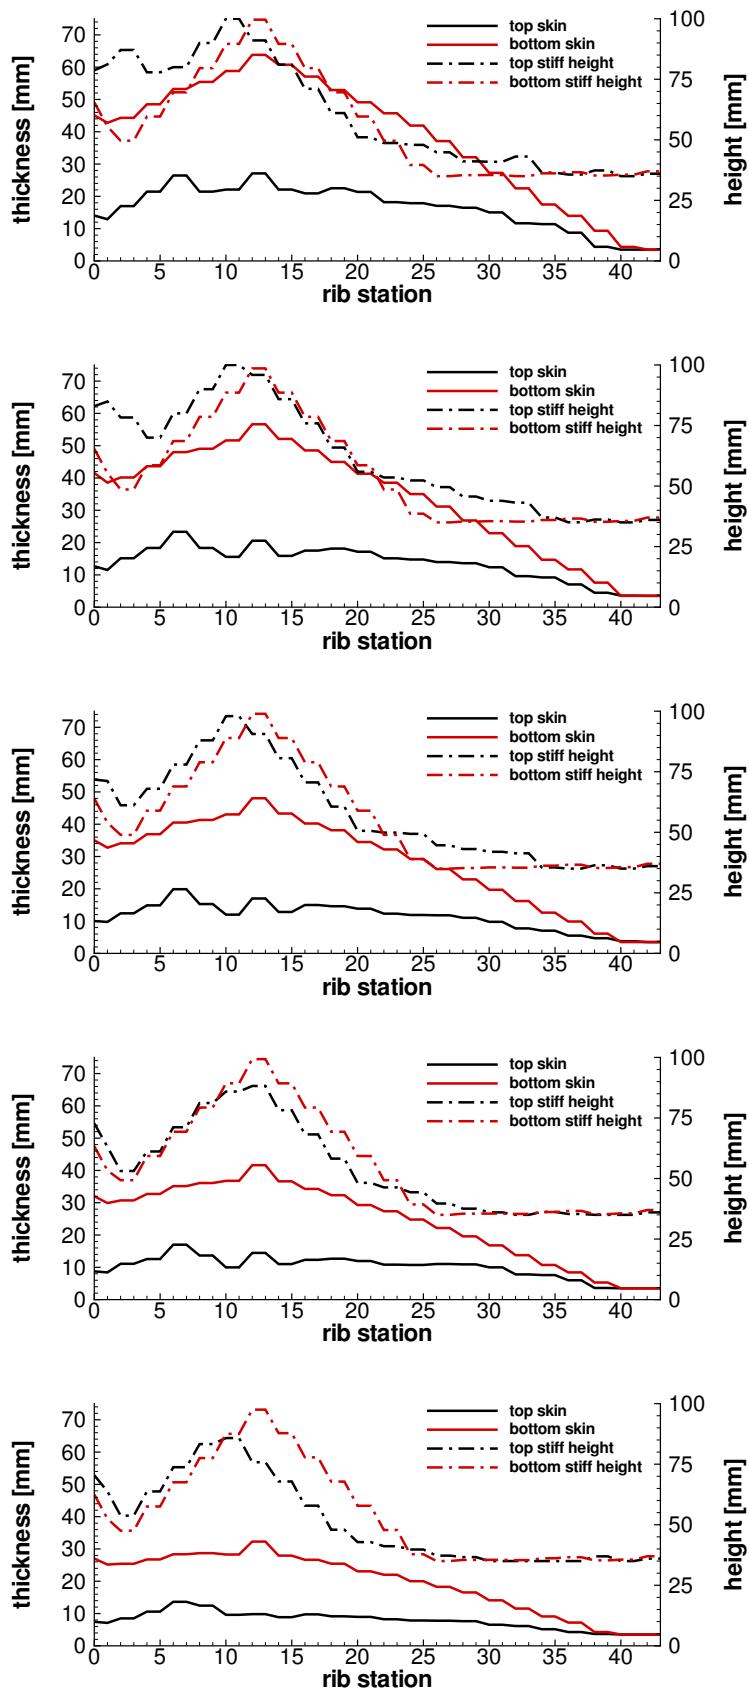


Figure 23: The skin thickness and stiffener height for the composite wings for $\beta = 1, 0.875, 0.75, 0.5,$ and 0 .

- [2] Jones, R. T. and Lasinski, T. A., "Effect of winglets on the induced drag of ideal wing shapes," Tech. Rep. TM-81230, NASA, September 1980.
- [3] Ning, S. A. and Kroo, I., "Multidisciplinary considerations in the design of wings and wing tip devices," *Journal of Aircraft*, Vol. 47, No. 2, 2010, pp. 534–543. doi:[10.2514/1.41833](https://doi.org/10.2514/1.41833).
- [4] Haftka, R. T., "Optimization of flexible wing structures subject to strength and induced drag constraints," *AIAA Journal*, Vol. 15, No. 8, August 1977, pp. 1101–1106.
- [5] Jansen, P., Perez, R. E., and Martins, J. R. R. A., "Aerostructural optimization of nonplanar lifting surfaces," *Journal of Aircraft*, Vol. 47, No. 5, 2010, pp. 1491–1503. doi:[10.2514/1.44727](https://doi.org/10.2514/1.44727).
- [6] Liu, B., Haftka, R. T., and Akgun, M. A., "Two-level composite wing structural optimization using response surfaces," *Structural and Multidisciplinary Optimization*, Vol. 20, 2000, pp. 87–96. doi:[10.1007/s001580050140](https://doi.org/10.1007/s001580050140).
- [7] Liu, B. and Haftka, R. T., "Single-level composite wing optimization based on flexural lamination parameters," *Structural and Multidisciplinary Optimization*, Vol. 26, 2004, pp. 111–120. doi:[10.1007/s00158-003-0315-6](https://doi.org/10.1007/s00158-003-0315-6).
- [8] Haftka, R. T. and Watson, L. T., "Multidisciplinary Design Optimization with Quasiseparable Subsystems," *Optimization and Engineering*, Vol. 6, 2005, pp. 9–20. doi:[10.1023/B:OPTE.0000048534.58121.93](https://doi.org/10.1023/B:OPTE.0000048534.58121.93).
- [9] Liu, B., Haftka, R. T., and Watson, L. T., "Global-local structural optimization using response surfaces of local optimization margins," *Structural and Multidisciplinary Optimization*, Vol. 27, 2004, pp. 352–359. doi:[10.1007/s00158-004-0393-0](https://doi.org/10.1007/s00158-004-0393-0).
- [10] Martins, J. R. R. A., Alonso, J. J., and Reuther, J. J., "High-fidelity aerostructural design optimization of a supersonic business jet," *Journal of Aircraft*, Vol. 41, No. 3, 2004, pp. 523–530. doi:[10.2514/1.11478](https://doi.org/10.2514/1.11478).
- [11] Martins, J. R. R. A., Alonso, J. J., and Reuther, J. J., "A coupled–adjoint sensitivity analysis method for high–fidelity aero–structural design," *Optimization and Engineering*, Vol. 6, 2005, pp. 33–62. doi:[10.1023/B:OPTE.0000048536.47956.62](https://doi.org/10.1023/B:OPTE.0000048536.47956.62).
- [12] Kenway, G. K. W., Kennedy, G. J., and Martins, J. R. R. A., "A Scalable Parallel Approach for High-Fidelity Aerostructural Analysis and Optimization," *53rd AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference*, Honolulu, HI, April 2012, AIAA-2012-1922.
- [13] Kennedy, G. J. and Martins, J. R. R. A., "Parallel solution methods for aerostructural analysis and design optimization," *Proceedings of the 13th AIAA/ISSMO Multidisciplinary Analysis Optimization Conference*, Fort Worth, TX, September 2010, AIAA 2010-9308.
- [14] Katz, J. and Plotkin, A., *Low-Speed Aerodynamics*, McGraw–Hill Inc., 1991.
- [15] Balay, S., Buschelman, K., Eijkhout, V., Gropp, W. D., Kaushik, D., Knepley, M. G., McInnes, L. C., Smith, B. F., and Zhang, H., "PETSc Users Manual," Tech. Rep. ANL-95/11 - Revision 2.1.5, Argonne National Laboratory, 2004.
- [16] Balay, S., Gropp, W. D., McInnes, L. C., and Smith, B. F., "Efficient Management of Parallelism in Object Oriented Numerical Software Libraries," *Modern Software Tools in Scientific Computing*, edited by E. Arge, A. M. Bruaset, and H. P. Langtangen, Birkhäuser Press, 1997, pp. 163–202.
- [17] Brown, S., "Displacement extrapolation for CFD+CSM aeroelastic analysis," AIAA Paper 97-1090, 1997.
- [18] Saad, Y. and Schultz, M. H., "GMRES: A generalized minimal residual algorithm for solving nonsymmetric linear systems," *SIAM Journal on Scientific and Statistical Computing*, Vol. 7, No. 3, 1986, pp. 856–869. doi:[10.1137/0907058](https://doi.org/10.1137/0907058).
- [19] Hopkins, E. J., "Charts for predicting turbulent skin friction from the van Driest method (II)," Tech. Rep. TN D-6945, NASA, October 1972.
- [20] Wakayama, S. and Kroo, I., "Subsonic wing planform design using multidisciplinary optimization," Vol. 32, No. 4, 1995, pp. 746–753. doi:[10.2514/3.46786](https://doi.org/10.2514/3.46786).
- [21] Valarezo, W. and Chin, V. D., *Journal of Aircraft*, Vol. 31, No. 1, 1994, pp. 103–109. doi:[10.2514/3.46461](https://doi.org/10.2514/3.46461).
- [22] Wrenn, G., "An indirect method for numerical optimization using the Kreisselmeier-Steinhauser function," NASA Technical Report CR-4220, 1989.
- [23] Kennedy, G. J. and Martins, J. R. R. A., "A Regularized Discrete Laminate Parametrization Technique with Applications to Wing-Box Design Optimization," *53rd AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference*, Honolulu, HI, April 2012, AIAA-2012-1519.
- [24] Haftka, R. T. and Walsh, J. L., "Stacking-sequence optimization for buckling of laminated plates by integer programming," *AIAA Journal*, Vol. 30, No. 3, March 1992, pp. 814–819.
- [25] Gill, P. E., Murray, W., and Saunders, M. A., "SNOPT: An SQP Algorithm for Large-Scale Constrained Optimization," *SIAM Review*, Vol. 47, No. 1, 2005, pp. 99–131.
- [26] Perez, R. E., Jansen, P. W., and Martins, J. R. R. A., "pyOpt: a Python-Based Object-Oriented Framework for Non-linear Constrained Optimization," *Structural and Multidisciplinary Optimization*, Vol. 45, No. 1, 2012, pp. 101–118. doi:[10.1007/s00158-011-0666-3](https://doi.org/10.1007/s00158-011-0666-3).

- [27] Baker, A. A., Dutton, S., and Kelly, D. W., *Composite Materials for Aircraft Structures*, AIAA Education Series, American Institute of Aeronautics and Astronautics, 2004.
- [28] Kenway, G. K. W., Kennedy, G. J., and Martins, J. R. R. A., "A CAD-free approach to high-fidelity aerostructural optimization," *13th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference, Forth Worth, TX*, AIAA-2010-9231, 2010.
- [29] Sederberg, T. W. and Parry, S. R., "Free-form deformation of solid geometric models," *SIGGRAPH Comput. Graph.*, Vol. 20, No. 4, 1986, pp. 151–160. doi:<http://doi.acm.org/10.1145/15886.15903>.
- [30] Hughes, P. C., *Spacecraft attitude dynamics*, Dover books on engineering, Dover Publications, 2004.
- [31] Ryan P. Henderson, J. R. R. A. M. and Perez, R. E., "Aircraft Conceptual Design for Optimal Environmental Performance," *The Aeronautical Journal*, Vol. 116, No. 1175, 2012, pp. 122.
- [32] Loken, C., Gruner, D., Groer, L., Peltier, R., Bunn, N., Craig, M., Henriques, T., Dempsey, J., Yu, C.-H., Chen, J., Dursi, L. J., Chong, J., Northrup, S., Pinto, J., Knecht, N., and Zon, R. V., "SciNet: Lessons Learned from Building a Power-efficient Top-20 System and Data Centre," *Journal of Physics: Conference Series*, Vol. 256, No. 1, 2010, pp. 012026.