High-fidelity aerostructural optimization considering buffet onset

Gaetan K. W. Kenway,[*]
Joaquim R. R. A. Martins[†]
University of Michigan, Department of Aerospace Engineering, Ann Arbor, MI

High fidelity computational modeling and optimization of complex engineering systems has the potential to allow engineers to produce more efficient designs with fewer unforeseen design modifications late in the design process. For transonic wing design, the simultaneous optimization of the aerodynamic shape and structural sizing yields significant fuel burn savings that can only be achieved with multidisciplinary optimization. Typically, optimizations attempt to improve the performance at a small number of representative operating conditions. However, the performance improvement may be obtained at the expense of reduced robustness or performance in other aspects of the design. In this paper, we investigate the effects of including operating conditions near the initial buffet onset boundary at both high C_L and high Mach number conditions. A series of aerostructural optimizations demonstrate the positive effect of including these off-design conditions on the buffet onset boundary of the optimized designs. A single point optimized design is able to reduce the fuel burn by 15.4% but with poor-off-design and buffet onset behavior. A multipoint optimization including conditions near the buffet onset, was able to reduce the fuel burn by 12.4%, improve the robustness near the on-design flight condition, and improve the buffet onset boundary.

I. Introduction

Numerical optimization has the potential to be a powerful tool that can complement more traditional design methodologies. Especially promising is optimization using high-fidelity, physics-based analysis methods such as Computational Fluid Dynamics (CFD) and Computational Structural Mechanics (CSM). By capturing the relevant physics of the underlying system, performance improvements predicted by the numerical analysis are more likely to be realized in the real world. Effective optimization algorithms, however, will invariably exploit any shortcomings in the numerical models and violate important constraints not included in the optimization problem formulation.

An example of how a missing physical effect in an optimization problem can be seen in previous optimization results by the authors [1]. In that work, multipoint optimization problems were solved with fuel burn and take-off gross weight (TOGW) objective functions. The TOGW objective favors designs with lower structural weight at the expense of higher drag. As a result, the optimized airfoil cross sections became excessively thick, approaching a 15% t/c near the Yehudi break, and require a rapid pressure recovery towards the trailing edge. These thick cross sections result in a substantial reduction of the wing box mass. However, since the solution of the Euler equations cannot predict flow separation, a RANS analysis of the optimized design showed worse performance than the baseline: The thick cross sections resulted in strong shocks that caused shock induced flow separation at the design condition. In this case, since the optimizer was missing the critical information necessary to prevent shock-induced flow separation, the optimized design violated the physical constraint.

A critical aspect of transonic wing design that is not usually considered in the optimization formulation is the behavior of the design near the buffet-onset boundary. As the Mach number or lift coefficient increases, shocks on the wing gradually increase in strength. Interactions with the boundary layer cause flow separation, and eventually this leads to unsteady separated flow and oscillation of the shock. This phenomena leads to an aerodynamically induced vibration of the airframe. Regulations stipulate that a 1.3 g margin must exist between the cruise condition and the initial buffet onset. By incorporating operating conditions near these margins in the optimization procedure, it is possible to allow the optimization to automatically determine the optimal cruise operating condition (C_L and altitude), while respecting the new buffet boundary of the optimized configuration.

There are a few reasons why we are particularly interested in the buffet onset boundary. First, with the large amount of design freedom that is possible with aerodynamic shape optimization, we would like to know the effects of subtle changes in the cross sectional shapes on the buffet boundary. Second, one of the key motivations for pursing

---

[*]Post Doctoral Research Fellow
[†]Associate Professor, AIAA Senior Member

high-fidelity aerostructural optimization is the possibility of including planform variables in the complete optimization procedure. This is a challenging proposition due to the far reaching consequences of planform changes on the rest of aircraft design. Changing the wing area in particular is particularly challenging: A change in wing area implies changing lift coefficients, which means the optimal altitude must be adjusted to maintain the best performance. Changing the altitude implies variable Reynolds numbers as well as variable engine thrust and thrust specific fuel consumption. Such flexibility, however, also allows the optimization algorithm to pick the best "nominal cruise" operating condition. In considering the buffet boundary, we want to ensure that the optimizer chooses a nominal operating condition that has a sufficient margin to buffet and it is, in fact, possible to operate with the performance predicted by the optimization algorithm.

This paper is organized as follows. First, we outline the main aspects of computational methods used this study. Then, we describe how we estimate the buffet boundary directly from CFD simulations. The next section describes a set of five optimization studies used to evaluate the effects of including buffet onset conditions, followed by concluding remarks.

II.   Computational Methods

The aerostructural analysis code used in this work is part of the MACH (MDO of Aircraft Configurations with High fidelity) framework. The tool is designed for aerostructural optimization studies of both conventional and unconventional aircraft configurations [2, 1].

1.   CFD Solver

The aerodynamic flow solver in MACH is SUMad, which solves the Reynolds averaged Navier–Stokes (RANS) equations in either steady, unsteady or time spectral modes [3, 4]. SUMad employs the finite-volume method on structured, body fitted, multi-block grids. The discretization scheme uses central fluxes with added artificial matrix dissipation and the Spalart–Allmaras (SA) turbulence model [5]. A fully coupled Newton–Krylov method is used to solve the mean flow and turbulence equations simultaneously. A discrete adjoint method is implemented using a combination of reverse mode automatic differentiation and analytic methods, for the efficient computation of the gradients of the functions of interest. More details on the CFD adjoint implementation are provided by Lyu *et al.* [6].

2.   Mesh Movement

The mesh movement algorithm used in this work is an analytic inverse distance method similar to the one described by Luke *et al.* [7]. Displacements for the full CFD volume mesh are a combination of all surface deformations weighted by the inverse of the distance to each surface node. The computational cost of a naive implementation of this method scales with the number of surface nodes. However, with a suitable fast spatial search algorithm and multipole-like expansion of the summation, the cost can be reduced to $\mathcal{O}\left(\log N\right)$. A careful implementation of the $\log N$ algorithm is fast enough and robust enough for use in aerostructural optimization. For a typical aerostructural analysis, the mesh movement scheme requires only 2–3% of the total solution time. The mesh movement procedure has also been differentiated using a combination of reverse mode automatic differentiation and analytic methods.

3.   Structural Solver

The MACH structural solver is the toolkit for the analysis of composite structures (TACS) presented by Kennedy and Martins [8, 9]. For the thin-shell problems typical of transport aircraft wing boxes, it is possible to have matrix condition numbers that exceed $\mathcal{O}\left(10^9\right)$. For this reason, we use a parallel direct method to solve the structural governing equations. TACS uses the adjoint method to efficiently compute the sensitivity of structural functions of interest with respect to structural and geometric design parameters.

4.   Aerostructural Solver

The aerostructural solver is responsible for the coupling of the aerodynamic and structural solvers, which are coupled using Python. This coupling includes determining how the aerodynamic loads are applied to the structural model and how the structural displacements are transferred to the aerodynamic model. The load and displacement transfer schemes closely follow the work of Brown [10]. In this approach, rigid links consisting of the shortest distance between the aerodynamic nodes and the structural surface are used to extrapolate the displacements from the structural surface to the CFD surface. The consistent force vector is determined by employing the method of virtual work, ensuring that the force transfer is conservative. The integration of the forces occurs on the aerodynamic mesh and is transmitted back through the rigid links to the structure.

In addition, the aerostructural solver is responsible for solving the coupled nonlinear system of equations and solving the adjoint equations. The nonlinear block Gauss–Seidel with Aitken acceleration method [11] is used for the nonlinear system, which has proven sufficiently robust for the range of flight conditions considered. The adjoint system uses the coupled Krylov method [2], which has a significant speed advantage over the traditional linear block Gauss–Seidel solver. The coupled aerostructural sensitivities are accurate to five or six digits, as all partial derivative terms are implemented with methods that are accurate to machine precision.

5. Engine Model

For the optimizations presented in Section V, the optimizer is able to modify the altitude of the various flight conditions. In addition to the aerodynamic consequences of varying the altitude, there are the thrust and thrust specific fuel consumption (TSFC) implications of altitude variation. We use a surrogate model of the engine performance map generated from a zero-dimensional engine model to estimate the thrust and thrust-specific fuel consumption for a 90 000 lb thrust class high-bypass ratio turbofan similar to a GE-90. An example of the variation in thrust specific fuel consumption as a function of Mach number and thrust setting, at an altitude of 37 000 ft, is shown in Figure 1.
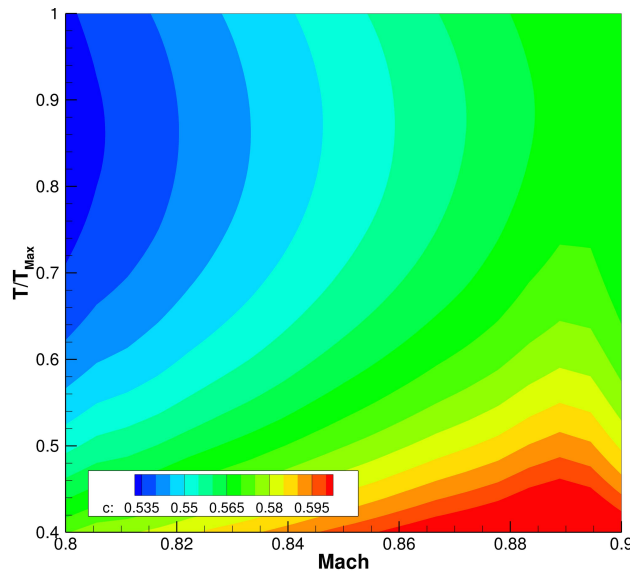

Figure 1. Thrust specific fuel consumption (c) contours for 37 000 ft given by the engine model.

6. Optimization Algorithm

The high computational cost of RANS-based aerostructural dictates that the optimization must perform a reasonably low number of function calls. We use SNOPT (sparse nonlinear optimizer) [12] through the Python interface py-Opt [13]. SNOPT is a gradient-based optimizer that implements a sequential quadratic programming method; it is capable of solving large-scale nonlinear optimization problems with thousands of constraints and design variables. SNOPT uses an augmented Lagrangian merit function, and the Hessian of the Lagrangian is approximated using a quasi-Newton method. We have successfully employed the SNOPT algorithm for a wide variety of aerodynamic and aerostructural optimization problems [1, 14, 15, 16, 17].

III. Buffet Onset Prediction

In transonic flow, at sufficiently high lift coefficients and Mach numbers, natural instabilities in the shock-boundary layer cause oscillations in the shock position and result in an unsteady flow field. Aerodynamic buffeting can be observed even on an idealized rigid wing. As aircraft approach the buffet boundary, these small aerodynamic oscillations grow in magnitude and induce increasingly large structural vibrations. Eventually, these oscillations may become large enough to cause structural failure or loss of control, especially if control surfaces are affected. Current regulations stipulate that commercial transport aircraft must maintain at least a 1.3 g margin from the cruise operating condition to the buffet onset boundary. This buffet margin may also be considered the maneuvering margin of the aircraft: A 1.3 g turn corresponds to an approximately 40° bank angle. Flight testing certification of new aircraft involves, among many

other things, mapping out the flight envelope up to and through the buffet boundary. In practice, for aircraft certification purposes, the buffet boundary is determined when the acceleration peak-to-peak amplitude reaches $n = 0.25$ at the pilot's seat. Such a definition is not particularly conducive for inclusion including in conceptual or preliminary design optimization. There are however, a multitude of other techniques that have been developed over the years that have been used to correlate data, typically wind tunnel data, with buffet predictions from the flight testing program. Some methods include RMS signal strength of wind tunnel model strain gauges, trailing edge pressure divergence, axial force break, pitching moment break, and lift curve break [18]. The last two methods can be readily applied to global data quantities obtained from CFD analysis. One particular way of implementing the lift curve break method is the $\Delta\alpha = 0.1$ method. Using this method, the linear portion of the lift curve is offset toward the right by a value of $0.1°$ and the intersection with the actual lift curve determines the buffet onset point. It is best explained graphically, as shown in Figure 2.
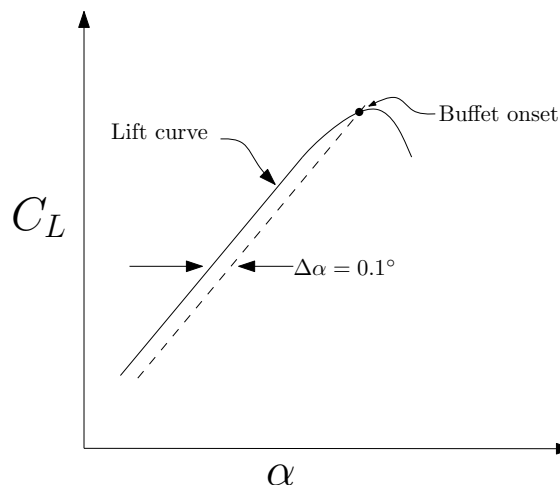

Figure 2. Estimating the buffet boundary with the $\Delta\alpha = 0.1$ method.

The physical mechanism that causes the loss of lift and the subsequent lowering of the lift curve slope is the initial appearance of shock induced flow separation. An example showing the typical progression at high angles of attack at fixed Mach number is given in the first two rows of Figure 3. The red surface shows the regions where the local surface flow is in the negative stream-wise (negative x-axis) direction. Alternatively, this is equivalent to the region where C_{f_x} is negative.

We have consistently observed that the appearance of this negative skin friction region correlates well with the initial buffet onset boundary as predicted by the $\Delta\alpha = 0.1$ method. This is perhaps not surprising, since the appearance of separated flow is the primary cause for the flattening of the lift curve slope. To that end, we have investigated using a "separation sensor" directly as the predictor for buffet onset. In its simplest form, the sensor is simply the integral of the area on the wing where the skin friction is negative. This function has a discontinuity as V_x crosses through zero. As a post-processing prediction tool, a discontinuity is acceptable; however, we would also like to use the separation sensor value directly in optimizations. For such an application it necessary to use a smooth Heaviside function instead. The exact form of separation sensor is given by:

$$\text{Sep} = \frac{1}{A_{\text{ref}}} \iint_S \frac{1.0}{1.0 + e^{-2k\chi}} dS$$
$$\chi = -\vec{V} \cdot \vec{V}_{\text{freestream}}$$

(1)

where k is a free parameter determining the sharpness of the transition. The values for $\vec{V}$ are taken from the state values at the cell center of the first off-wall cell, as the velocities are zero by definition for a no-slip wall condition. Both $\vec{V}$ and $V_{\text{freestream}}$ have been normalized to a unit magnitude, ensuring the dot product falls between -1 and 1. The specific smooth Heaviside function with $k = 10$ used in this work is shown in Figure 4.

The last row of Figure 3 shows the value of χ on the wing surface for three angles of attack at $M = 0.85$. The effect of the smooth Heaviside function is to smear out the sensor value around the separated flow region.

To demonstrate the utility of the separation sensor-based method for predicting buffet onset, we have computed the buffet boundary for the baseline CRM configuration at a flight altitude of 37,000 ft over the Mach range 0.8 to 0.9
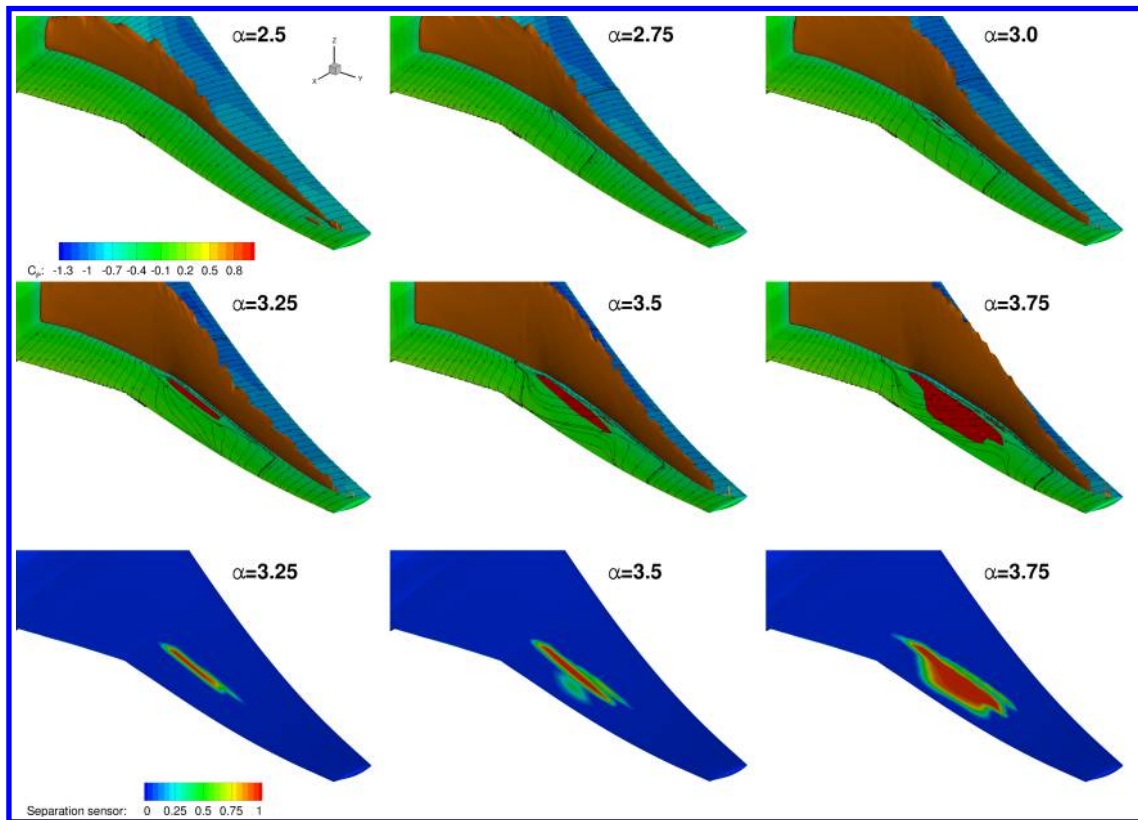
American Institute of Aeronautics and Astronautics

Figure 3. Progression of separated flow for the baseline CRM configuration at $M = 0.85$**. The red surface denotes flow in the negative streamwise direction. The orange isosurface is the Lovely–Haines shock sensor [19]. The last row shows the separation sensor integrand value from Eqn.** (1)**.**
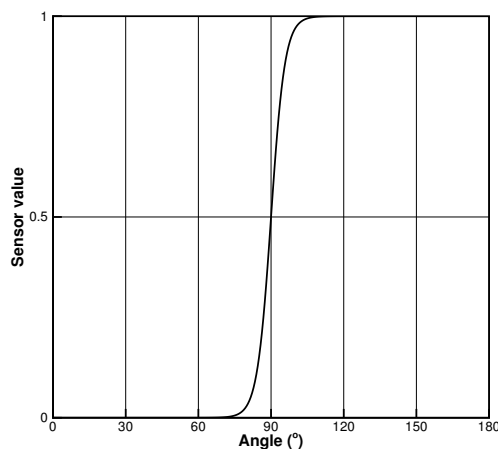

Figure 4. With $k = 10$**, the majority of the smooth variation occurs over** $\pm 15°$**.**

using the $\Delta \alpha = 0.1$ method as well as with the separation criteria. The separation criteria requires the specification of a cutoff value. For this work, we have determined this value by comparing the resulting buffet boundary with the $\Delta \alpha = 0.1$ method. The comparison with Sep $= 3.3\%$ is shown in Figure 5. Overall agreement between the two methods is good, with the overall shape of buffet boundary curve matching well.

Figure 6 demonstrates how the separation sensor approach is used to predict the buffet onset boundary. For each α sweep at a fixed Mach number, the intersection of the separation value with the specified cutoff value determines the

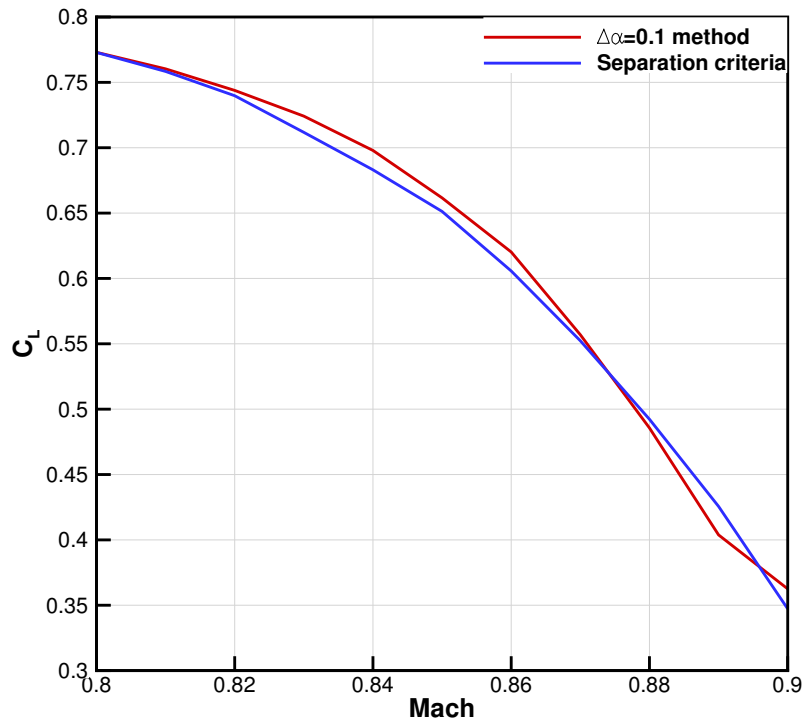American Institute of Aeronautics and Astronautics

Figure 5. Comparison of $\Delta\alpha = 0.1$ method and separation sensor method.

boundary. Note that the sensor value is nonzero even at low angles of attack. The reason for this is there is a small amount of reverse flow near the wing and tail leading edges due to the lack of leading edge extensions.
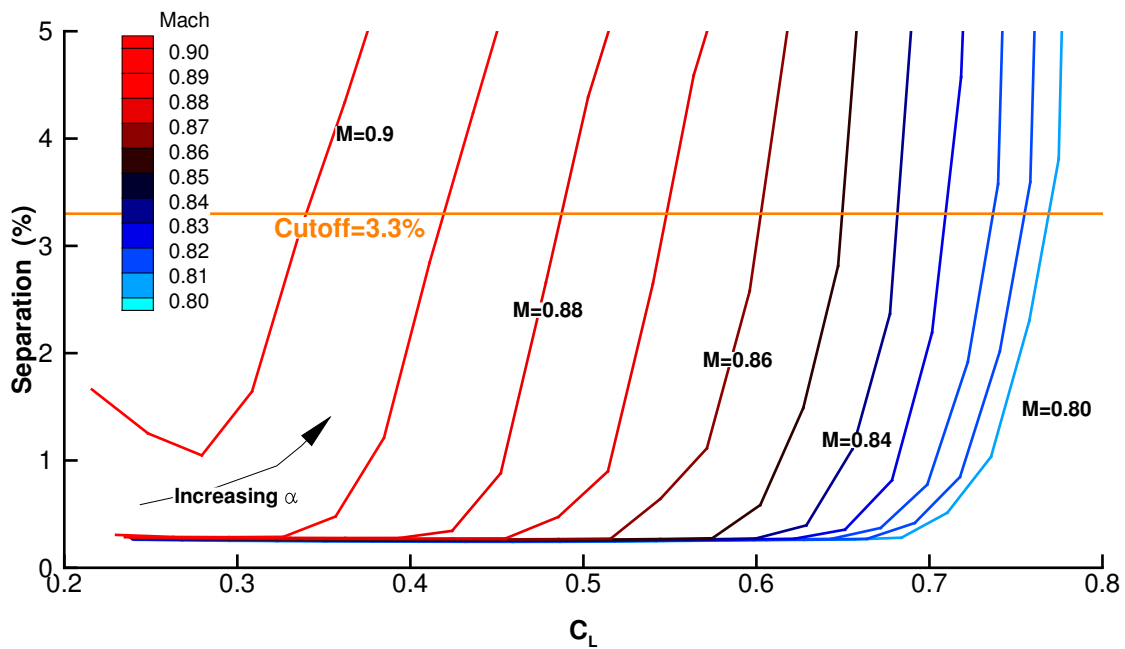

Figure 6. Separation sensor value for α sweeps at a range of Mach numbers. The cutoff value indicates the estimated buffet boundary.

American Institute of Aeronautics and Astronautics

IV. Optimization Problem Description

1. Geometry Overview

The geometry under consideration is based on the NASA Common Research Model (CRM) wing-body-tail configuration, used previously in the 4^{th} and 5^{th} drag prediction workshops as well as the Aerodynamic Design Optimization Discussion Group test cases [20, 21, 22, 23, 24] and previous aerostructural optimization work by the authors [1, 14].

The CRM geometry is only provided in a nominal 1 g configuration, which makes it initially unsuitable for aerostructural analysis. In previous work by the authors [14], we used an inverse design procedure to determine the jig shape that deforms to the flying shape at a specified nominal operating condition of $M = 0.85$, $C_L = 0.5$, altitude$= 37\,000$ ft. This undeformed geometry, which we refer to as the uCRM is the starting initial geometry for this study. An overview of the initial uCRM configuration is given in Figure 7.
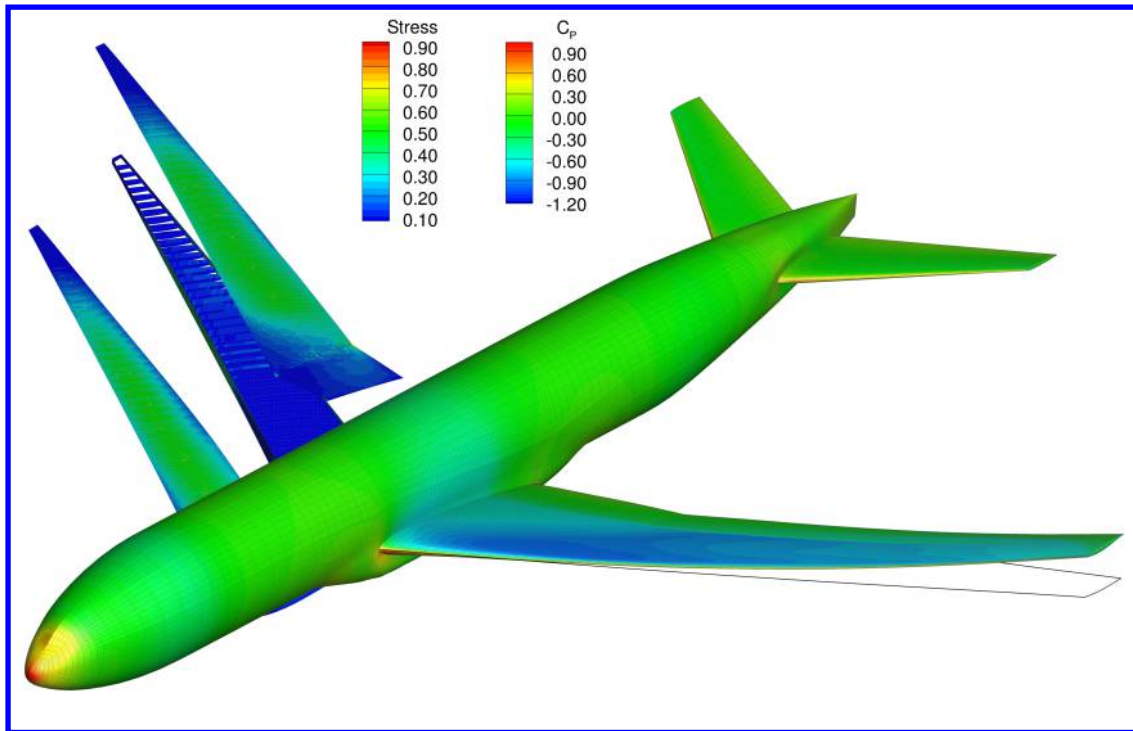

Figure 7. uCRM aerostructural model. The wingbox structural model is shown on the left and the aerodynamic model on the right. The black outline denotes the jig shape. The CFD surface mesh resolution is also visible.

The overall size of the uCRM configuration is similar to a Boeing 777-200ER. The additional information required to define the aerostructural optimization problem is obtained using publicly available data [25] (see Table 1).

2. Objective Function and Operating Conditions

The objective for each optimization is to reduce the following composite function that involves both fuel burn and takeoff gross weight (TOGW):

$$\text{Obj} = \sum_{i}^{N} \mathcal{W}_i \left(\beta \times \text{Fuel Burn} + (1.0 - \beta) \, \text{TOGW} \right). \tag{2}$$

The parameter β is used to select a blending of the fuel burn and TOGW. Extreme β values of 0.0 and 1.0 yield the TOGW and fuel burn objectives, respectively. The real objective function for a transport aircraft lies somewhere between these two extremes and depends on how the airline uses the aircraft, fuel price, and other factors.

For each flight condition described below, we compute the fuel burn and TOGW as if the entire mission was completed in a cruise-climb fashion at the L/D ratio evaluated at that analysis point. For simplicity, we ignore the fuel burn associated with the taxi, take-off climb, and descent phases. Using the Breguet range equation applied to the

American Institute of Aeronautics and Astronautics

Table 1. uCRM specifications

Parameter	Value	Units
Cruise Mach number	0.85	-
Cruise lift coefficient	0.5	-
Cruise altitude	37 000	ft
Span	58.8	m
Aspect ratio	9.0	-
Reference wing area	383.7	m^2
Sweep (leading edge)	37.4	°
Maximum take-off weight (MTOW)	297 500	kg
Operational empty weight	143 000	kg
Design range	7 250	nm
Design payload	34 000	kg
Reserve fuel	15 000	kg
Fixed weight	101 468	kg

full design range of 7 250 nm, we may compute the TOGW and fuel burn objectives according to:

$$\text{TOGW} = \text{LGW} e^{\left(\frac{R\,\text{TSFC}}{V(L/D)}\right)}$$

$$\text{FB} = \text{TOGW} - \text{LGW}$$

(3)

where FB is the fuel burn, LGW is the aircraft landing weight, R is the design mission range, TSFC is the thrust-specific fuel consumption, V is the cruise speed, and L/D is the lift-to-drag ratio. Since we are only modeling the wing, body and tail, an additional 30 counts of drag are added to the CFD computed drag coefficient to account for the vertical stabilizer, nacelles and pylons. The landing weight is computed using the following formula:

$$\text{LGW} = 1.25 \times W + \text{Fixed Weight} + \text{Payload} + \text{Reserve Fuel Weight} + \text{Secondary Wing Weight} \frac{A}{A_{\text{ref}}} \quad (4)$$

where W is the weight computed by the structural finite-element model. The factor of 1.25 accounts for additional weight associated with fasteners, overlaps, and other components not modeled in the idealized wingbox. The Secondary Wing Weight is used to account for the changes in wing mass that scale with the exposed area, A. This is important, as we expect the wing mass to increase as the area increases, independent of the primary structural weight. The Secondary Wing Weight is set at 4 000 kg per wing.

Since the goal of this work is to evaluate the effect of including buffet onset conditions in the design optimization, we formulate a series of four optimizations that allow us to examine the effects of these conditions on both single and multipoint designs. It is well known that single point aerodynamic optimizations result in excellent on-design performance at the expense of poor off-design performance. Multipoint optimizations, on the other hand, use the combined performance at a variety of operating conditions in an attempt to obtain a more robust design. One question we would like to answer is if an extensive multipoint optimization with all operating conditions near the nominal operating condition is sufficient to ensure robust buffet onset behavior. Table 2 describes the operating conditions for each of the four optimizations.

Case 1 is a simple single point optimization that is included to serve as a reference. Case 2 is an equally weighted seven-point optimization. The design points are selected to be representative conditions that aircraft would experience during normal operations. Cases 3 and 4 mirror Cases 1 and 2, but include two off-design points that are intended to limit poor aerodynamic behavior near the buffet boundary. A graphical representation of the operating conditions is shown in Figure 8.

The initial basic operating condition for Case 1 is $M = 0.85$, $C_L = 0.5$, altitude=34 000 ft. The lift at this operating condition corresponds to a point in the flight where 1/2 of the total mission fuel has been burned. The next four operating conditions, Points 2 through 5, are evaluated at the same physical lift (fuel fraction of 0.5). Cases 2 and 3 vary the Mach number by ± 0.01, and the C_L is modified accordingly. Points 4 and 5 vary the altitude by ± 2000 ft, which yields larger variations in the lift coefficient. The final two normal operating conditions—Points 6 and 7—vary the fuel fraction from near empty (10%) to nearly full (90%). The main reason for including these points, is to capture the effects due to higher or lower wing bending and the accompanying twist variations. For these cases the altitudes have been adjusted so that the C_L is close to the nominal value. Note that the weights given in Figure 8 are only

American Institute of Aeronautics and Astronautics

Table 2. Operating conditions for each optimization.

Case	Point	Weights ($\mathcal{W}_i$)	Mach	Fuel Fraction	Altitude (ft)	Load Factor
1	1	1.0	0.85	0.5	Design	1.0
	1	1/7	0.85	0.5	Design	1.0
	2	1/7	0.84	0.5	Design	1.0
	3	1/7	0.86	0.5	Design	1.0
2	4	1/7	0.85	0.5	Design -2000	1.0
	5	1/7	0.85	0.5	Design $+2000$	1.0
	6	1/7	0.85	0.1	Design $+4000$	1.0
	7	1/7	0.85	0.9	Design -3000	1.0
	1	4/5	0.85	0.5	Design	1.0
3	2	1/10	0.85	0.5	Design	1.3
	3	1/10	0.89	0.5	Design	1.0
	1	8/70	0.85	0.5	Design	1.0
	2	8/70	0.84	0.5	Design	1.0
	3	8/70	0.86	0.5	Design	1.0
	4	8/70	0.85	0.5	Design -2000	1.0
4	5	8/70	0.85	0.5	Design $+2000$	1.0
	6	8/70	0.85	0.1	Design $+4000$	1.0
	7	8/70	0.85	0.9	Design -3000	1.0
	8	1/10	0.85	0.5	Design	1.3
	9	1/10	0.89	0.5	Design	1.0

notional and actually vary during the optimization as the fuel burn changes. The specified fuel fraction is the defining characteristic that determines the overall aircraft weight.

The final two operating conditions for Cases 3 and 4 are the buffet onset conditions. For this optimization study, we have chosen to include these two conditions directly in the objective function. The first buffet condition is chosen to be 1.3 g higher than the nominal design point (Point 1). The second buffet point, is at the maximum operating Mach (M_{MO}) condition, which we have assumed to be $M = 0.89$ for the CRM configuration. The lift at this condition is the same as the nominal flight condition and thus the increased Mach number results in a lower coefficient of lift.

It is also necessary to determine how to include the buffet operating conditions in the objective function. From preliminary aerodynamic optimizations, it was clear that for an optimization such as Case 3, equally weighting the nominal flight conditions with the two off-design conditions was undesirable. With some trial and error, we determined that a weighting of $4/5$ for the on-design condition and $1/10$ for each buffet condition resulted in fully attached or nearly fully attached flow at the off-design conditions, while only mildly impacting the on-design performance. For the multipoint optimization, Case 4, the buffet weights are kept constant at $1/10$, and the single point weighting of $4/5$ was evenly distributed to the now 7 normal operating conditions. This actually results in the normal operating conditions being weighted only slightly higher than the buffet conditions.

In addition to the cruise and buffet operating conditions, we must also define several maneuver loading conditions that will produce representative loads to size the wing structure. The three maneuver load conditions are as follows:

2.5 g symmetric pull up : For transport aircraft this is one of the most critical sizing conditions. This analysis is performed at TOGW, $M = 0.6415$, sea-level altitude with full fuel. This Mach number is chosen to correspond to an assumed V_D speed of 425 knots.

-1.0 g symmetric push over : The primary effect of this maneuver condition is the sizing of the lower wing surface dictated by local buckling constraints. The analysis is performed at the same conditions as the 2.5 g condition.

Quasi-steady gust load : Previous results have shown that for highly flexible aerostructurally optimized wings, it is possible to have a structural constraints that are critical at a 1 g cruise condition [14]. Therefore, the purpose of this condition is to ensure the operating stresses at normal operating conditions are not too high. It can also be considered as a surrogate for a gust load that may be expected to load the wing sufficiently quickly that the passive aeroelastic tailoring cannot redistribute the load inboard fast enough. This maneuver condition is analyzed near the Mach cross-over altitude of 28 000 ft, $M = 0.85$, TOGW and full fuel load.
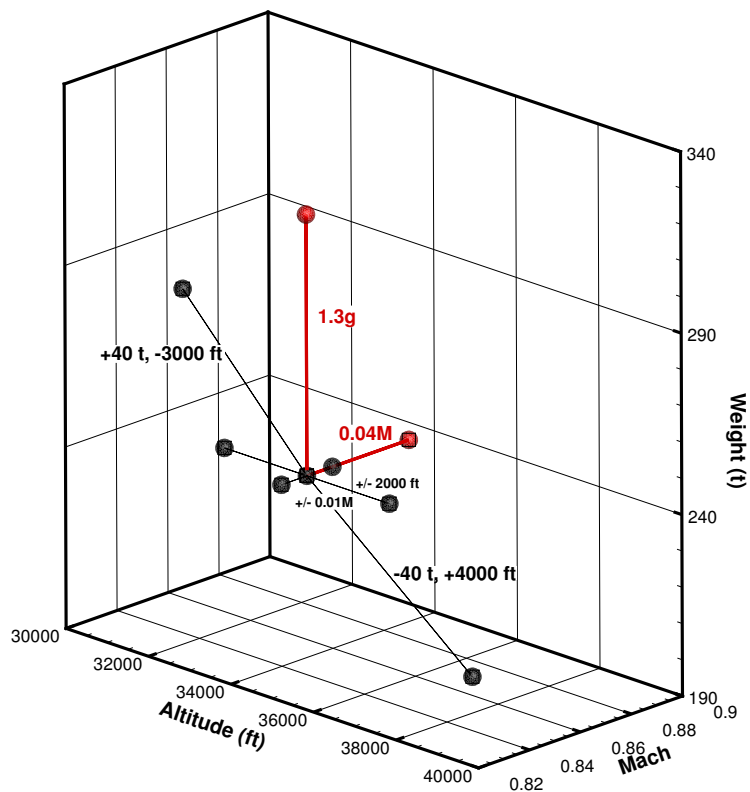
American Institute of Aeronautics and Astronautics

Figure 8. Flight operating conditions considered in the optimizations; only Case 4 uses all of these operating conditions.

The wingbox structure is assumed to be constructed of a 7000 series aluminum alloy with a limiting stress of 420 Mpa. The 2.5 g and −1.0 g conditions incorporate an additional 1.5 safety factor as required by regulations. The third flight condition utilizes a 2.67 safety factor to ensure a sufficient structural margin for a gust encountered near the cruising altitude.

3. Design Variables and Constraints

Using the coupled adjoint method, the cost of computing gradients of functions is nearly independent of the number of design variables [2] allowing us to fully exploit the benefits of high-fidelity aerostructural optimization. There are several groups of design variables: geometric design variables, aerodynamic variables, engine model variables, and structural variables.

We have selected a compact set of variables that are able to efficiently manipulate the wing planform: span, sweep, and chord. When combined, the span and chord variables control the wing area and aspect ratio. Since the chord variable is applied evenly to the entire span, the taper ratio of the initial configuration remains constant. The planform variables yield tremendous design freedom typically not found in high-fidelity aerodynamic only optimization. In particular, the optimization process is free to explore the trade-off between increasing the wingspan to reduce the span-loading and thus induced drag at the expense of a potentially heavier structure. Likewise, the sweep angle may be modified in response to changes in the airfoil thickness distributions as well as accounting for the change in structural mass. The remaining geometric variables are 7 twist variables, a tail rotation angle for each flight condition, and 192 shape variables controlling the detailed cross sections.

The aerodynamic variables include the angle of attack at each of the cruise and maneuver conditions and the nominal cruise altitude. This single altitude variable is used to drive the operating altitude for all flight conditions. There is an additional target variable X_{ref}, which modifies the position about which the moments are taken. This is then constrained to be the same as the 1/4 chord MAC position of the trapezoidal wing.

Each of the normal operating conditions have a throttle variable. This variable determines the thrust requested from the engine, which is necessary to meet the $D = T$ constraint.

The remainder of the design variables are used for the parametrization of the structural wingbox. The first four

American Institute of Aeronautics and Astronautics

variables are the stiffener pitch of the upper skin, lower skin, leading edge spar and trailing edge spar. We assume that the stiffener pitch is thus constant across each component. The panel-based smeared stiffness approach results in three design variables for the skin or spar in each rib bay: panel thickness, stiffener thickness, and stiffener height. An additional variable, the panel length, is used to simplify panel buckling computations and is constrained to match the physical panel length (which changes as a function of the geometric design variables) through an equal number of nonlinear constraints. Altogether, there are 972 design optimization variables in Case 4. The other three optimizations have slightly fewer variables and constraints due to the lower number of operating conditions. A graphical depiction of geometric and structural design variables is shown in Figure 9.
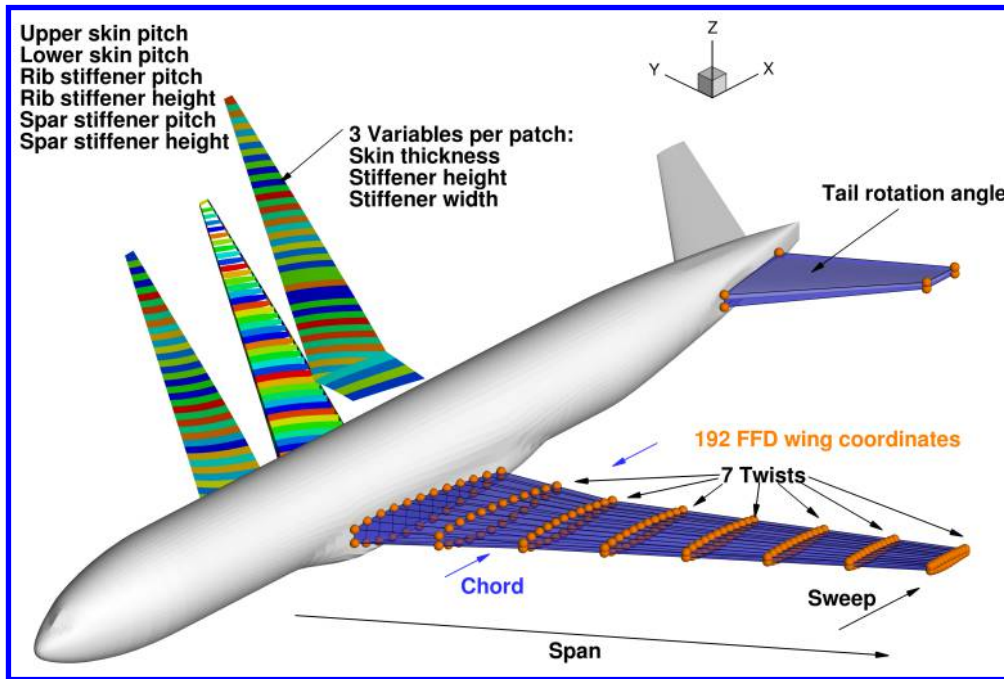

Figure 9. Geometric and structural optimization variables.

Even with high fidelity analyses, the optimization problem requires many constraints. Each of the cruise and maneuver operating conditions require that lift equals weight and that the configuration is trimmed ($C_{M_y} = 0.0$). Each of the normal operating conditions have two additional constraints related to the engine model. The first is that the thrust must equal the drag, which is met by the engine throttle setting variable. This is important, as the TSFC is a function of throttle setting. The second is an excess thrust constraint, which ensures sufficient climb margin (or equivalently thrust to drag ratio) at the normal operating conditions.

A variety of geometric constraints are used to ensure a realistic design is produced. The leading edge radius constraint ensures no reduction in leading edge radius, potentially causing reduction in the low-speed $C_{L_{max}}$ performance. Previous aerodynamic optimizations have shown that single point optimized designs without this constraint produce sharp leading edges in transonic flow [16]. The trailing edge constraints are used to ensure a manufacturable trailing edge that is not too thin.

The fuel volume constraint ensures that there is sufficient internal wing volume to store the required fuel. Note the coupled nature of this constraint: if the optimizer lowers the fuel burn, the wing volume is allowed to decrease. The 2.5 g maneuver condition uses four KS stress constraint aggregation functions: one for the upper wing surface, a second for the lower wing surface, one for the spars, and a final one for the ribs. Three buckling constraints are used for the upper skin, spars, and ribs respectively. For the -1.0 g condition, three buckling constraints are used: one for the lower skin, one for the spars and a final for the ribs. Finally, for the quasi-steady gust load condition, there are three yield stress constraints and four buckling constraints. We use several hundred adjacency constraints that ensure there are no large changes in the properties of the skin or stiffener between adjacent panels. Finally, there are 16 constraints for the FFD airfoil shape variables that keep the leading and trailing edges from moving in the z direction. This is required to eliminate a "shearing twist" design mode, and allows the pure rotation (nonlinear) twist to modify the overall wing twist. In total, there are 961 optimization constraints, for the case 4 optimization. The full optimization problem for Case 4 is summarized below.

	Function/variable	Description	Quantity
minimize	$\sum_i^N \mathcal{W}_i \left(\beta \times \text{Fuel Burn} + (1.0 - \beta)\, \text{TOGW} \right)$		
with respect to	x_{span}	Wing span	1
	x_{sweep}	Wing sweep	1
	x_{chord}	Wing chord	1
	x_{twist}	Wing twist	7
	x_{airfoil}	FFD control points	192
	x_{alpha_i}	Angle of attack at each flight condition	12
	x_{η_i}	Tail rotation angle at each flight condition	12
	x_{throttle_i}	Throttle setting for each cruise flight condition	7
	x_{altitude}	Cruise altitude	1
	X_{CG}	CG position	1
	$x_{\text{skin pitch}}$	Upper/lower stiffener pitch	2
	$x_{\text{spar pitch}}$	LE/TE Spar stiffener pitch	2
	x_{ribs}	Rib thickness	45
	$x_{\text{panel thick}}$	Panel thickness skins/spars	172
	$x_{\text{stiff thick}}$	Panel stiffener thickness skins/spars	172
	$x_{\text{stiff height}}$	Panel stiffener height skins/spars	172
	$x_{\text{panel length}}$	Panel length skin/spars	172
		Total design variables	**972**
subject to	$L = n_i W$	Lift constraint	12
	$C_{M_{y_i}} = 0.0$	Trim constraint	12
	$T = D$	Thrust constraint	7
	$1.08D - T_{\max} < 0$	Excess thrust constraint	7
	$t_{\text{LE}}/t_{\text{LE}_{\text{Init}}} \geq 1.0$	Leading edge radius	20
	$t_{\text{TE}}/t_{\text{TE}_{\text{Init}}} \geq 1.0$	Trailing edge thickness	20
	$\mathcal{V}_{\text{wing}} > \mathcal{V}_{\text{fuel}}$	Minimum fuel volume	1
	$x_{\text{CG}} - 1/4 MAC = 0$	CG location at 1/4 chord MAC	1
	$L_{\text{panel}} - x_{\text{panel length}} = 0$	Target panel length	172
	$\text{KS}_{\text{stress}} \leq 1$	2.5 g Yield stress	4
	$\text{KS}_{\text{buckling}} \leq 1$	2.5 g Buckling	3
	$\text{KS}_{\text{buckling}} \leq 1$	-1.0 g Buckling	3
	$\text{KS}_{\text{buckling}} \leq 1$	1.78 g Yield stress	3
	$\text{KS}_{\text{buckling}} \leq 1$	1.78 g Buckling	4
	$\left\| x_{\text{panel thick}_i} - x_{\text{panel thick}_{i+1}} \right\| \leq 0.005$	Skin thickness adjacency	168
	$\left\| x_{\text{stiff thick}_i} - x_{\text{stiff thick}_{i+1}} \right\| \leq 0.005$	Stiffener thickness adjacency	168
	$\left\| x_{\text{stiff height}_i} - x_{\text{stiff height}_{i+1}} \right\| \leq 0.005$	Stiffener height adjacency	168
	$x_{\text{stiff thick}} - x_{\text{panel thick}} < 0.005$	Maximum stiffener-skin difference	172
	$\Delta z_{\text{TE,upper}} = -\Delta z_{\text{TE,lower}}$	Fixed trailing edge	8
	$\Delta z_{\text{LE,upper}} = -\Delta z_{\text{LE,lower}}$	Fixed leading edge	8
		Total constraints	**961**

V.   Results

In this section we provide an overview of the most significant results from each optimization. The four optimization problems described in the previous section all minimize fuel burn as the objective, i.e., $\beta = 1.0$ in Eqn. (2). We also introduce one additional optimization, Case 4b, which is identical to Case 4, but with TOGW as the objective function (i.e., $\beta = 0.0$). This provides an opportunity to see the effect of the objective choice on the optimized designs.

First we look at the overall trends of several key properties for each of the five optimizations shown in Figure 10. It is immediately clear that the choice of objective has a large overriding impact on the overall design. For almost all parameters the overall trend of the TOGW optimization is opposite of the fuel burn objective optimizations (Cases 1 through 4). This dramatic effect of the objective function on the optimized design confirms the same overall trends

described previously by the authors using an Euler-based aerostructural optimization [2]. The overall trends of the fuel burn minimization show significant increases in the wing area, aspect ratio and lift to drag ratio. As a consequence of the increased wing area, the design altitude is increased to its specified limit of 39 000 ft. As a result of the increased aspect ratio and span, the fuel burn optimizations significantly increase the wingbox structural mass from approximately 41 000 kg up to approximately 55 000 kg for Case 4. It is interesting to note that for Cases 3 and 4, the TOGW (based on the design condition) has actually increased. This can be explained by the fact that while the fuel burn has decreased (as expected since this is the objective function), the aircraft empty weight has increased more than the fuel burn has decreased. The opposite effect can be seen with Case 4b: The fuel burn has actually increased, but the TOGW has been reduced due to a larger reduction in the wingbox structural weight.

If we look closer at the variations between the fuel burn optimizations (Cases 1 through 4), we can see the impact of increasing number of off-design flight conditions. The lowest fuel burn (objective function) is obtained for the single point optimization (Case 1) as is expected. Case 1 also has the lowest wing area, the lowest wing mass, lowest sweep angle and the highest aspect ratio. The next lowest fuel burn is Case 2, and correspondingly has the second lowest sweep, wing wing and wingbox max. Case 3 and 4, both of which include the off-design buffet conditions have very similar on-design performance. These cases also have the highest wing sweep, although it should be noted, it is still 1 degree lower than the baseline design. It is clear that the additional buffet operating conditions have a noticeable impact on the key variables. Interestingly, all the fuel burn optimizations yield very similar L/D values, ranging from 23.60 to 23.75. This indicates the variation in the fuel burn between the cases is primarily due to the lower wingbox mass of Case 1 and 2.

Next, we take a closer look at the main results of each optimization. Figures 11 through 15 present a broad overview of the key aspects of the optimized aerodynamic and structural designs. The figures show the C_p and planform comparison versus the initial design (upper left), the equivalent thickness distribution, stress and buckling KS failure criteria (upper right), a comparison of initial and optimized lift distributions, twist distributions, and thickness to chord ratio (t/c) (lower left), and four airfoils with corresponding C_p distributions (lower right). There is also a visual comparison of the deflection of the -1.0 g maneuver condition, the initial jig shape, the nominal cruise condition, and the 2.5 g maneuver condition.
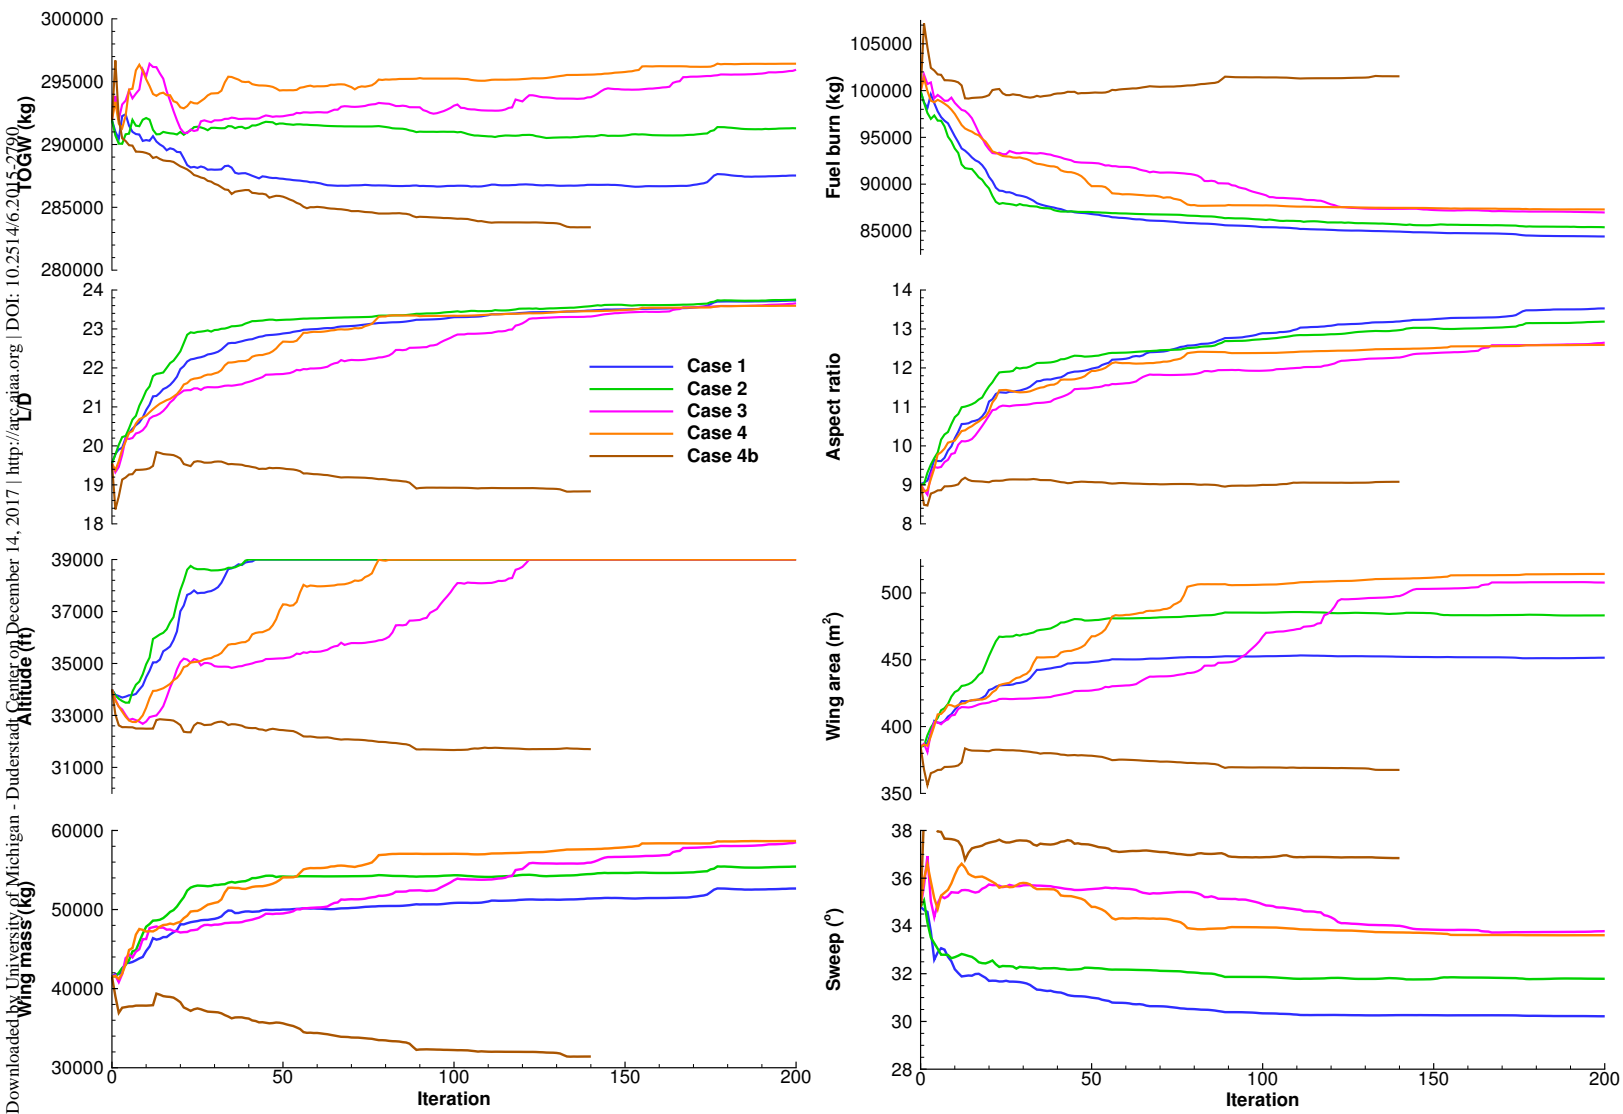
American Institute of Aeronautics and Astronautics

Figure 10. Evolution of key parameters for each optimization.
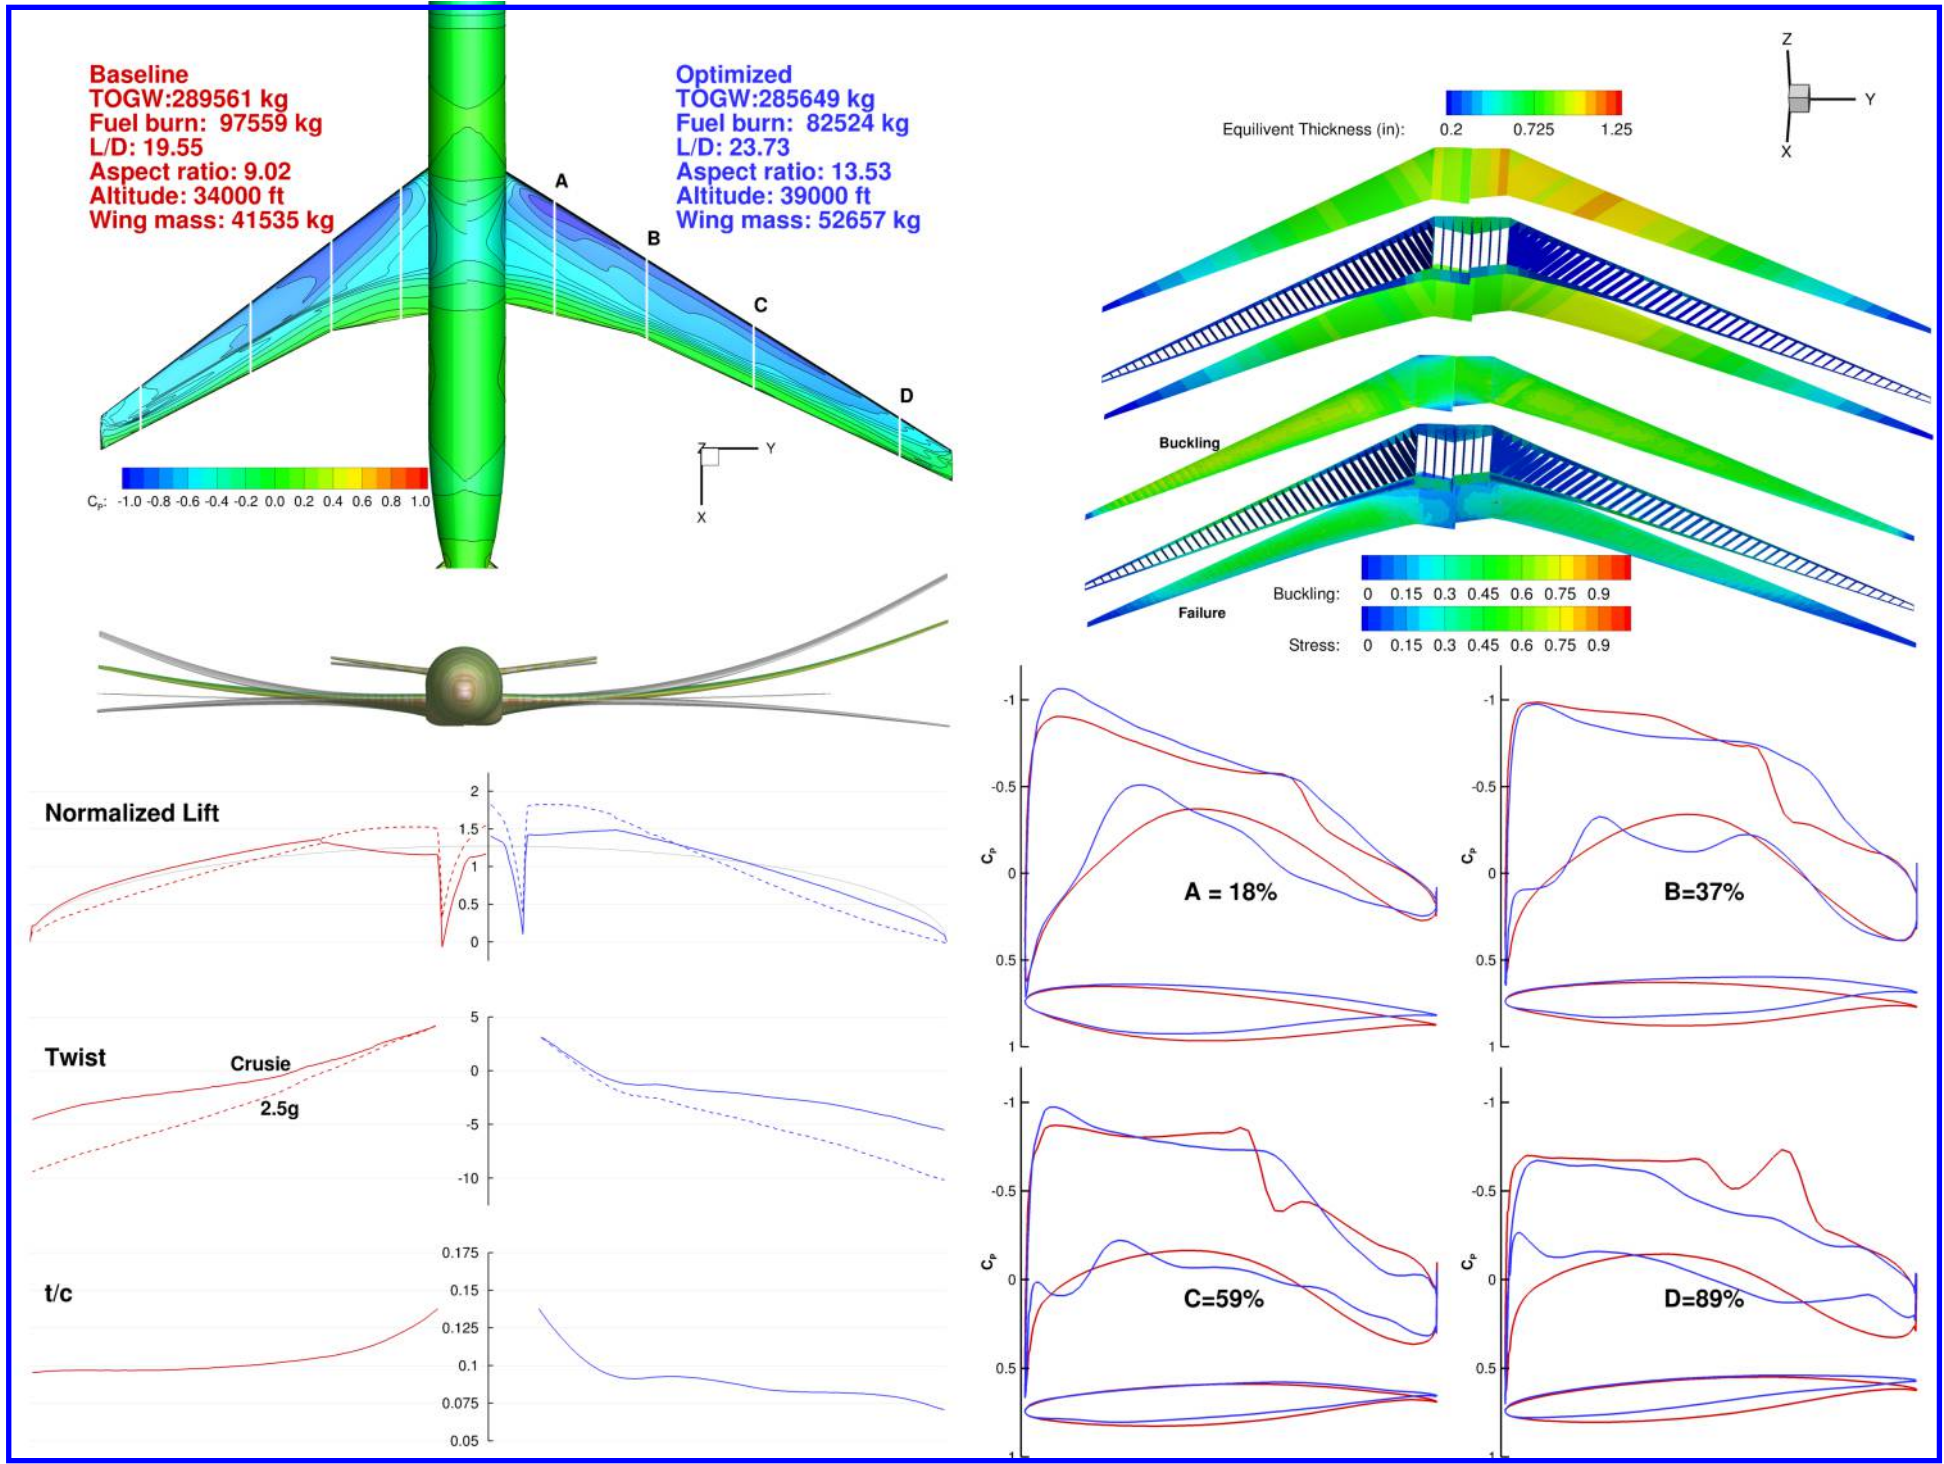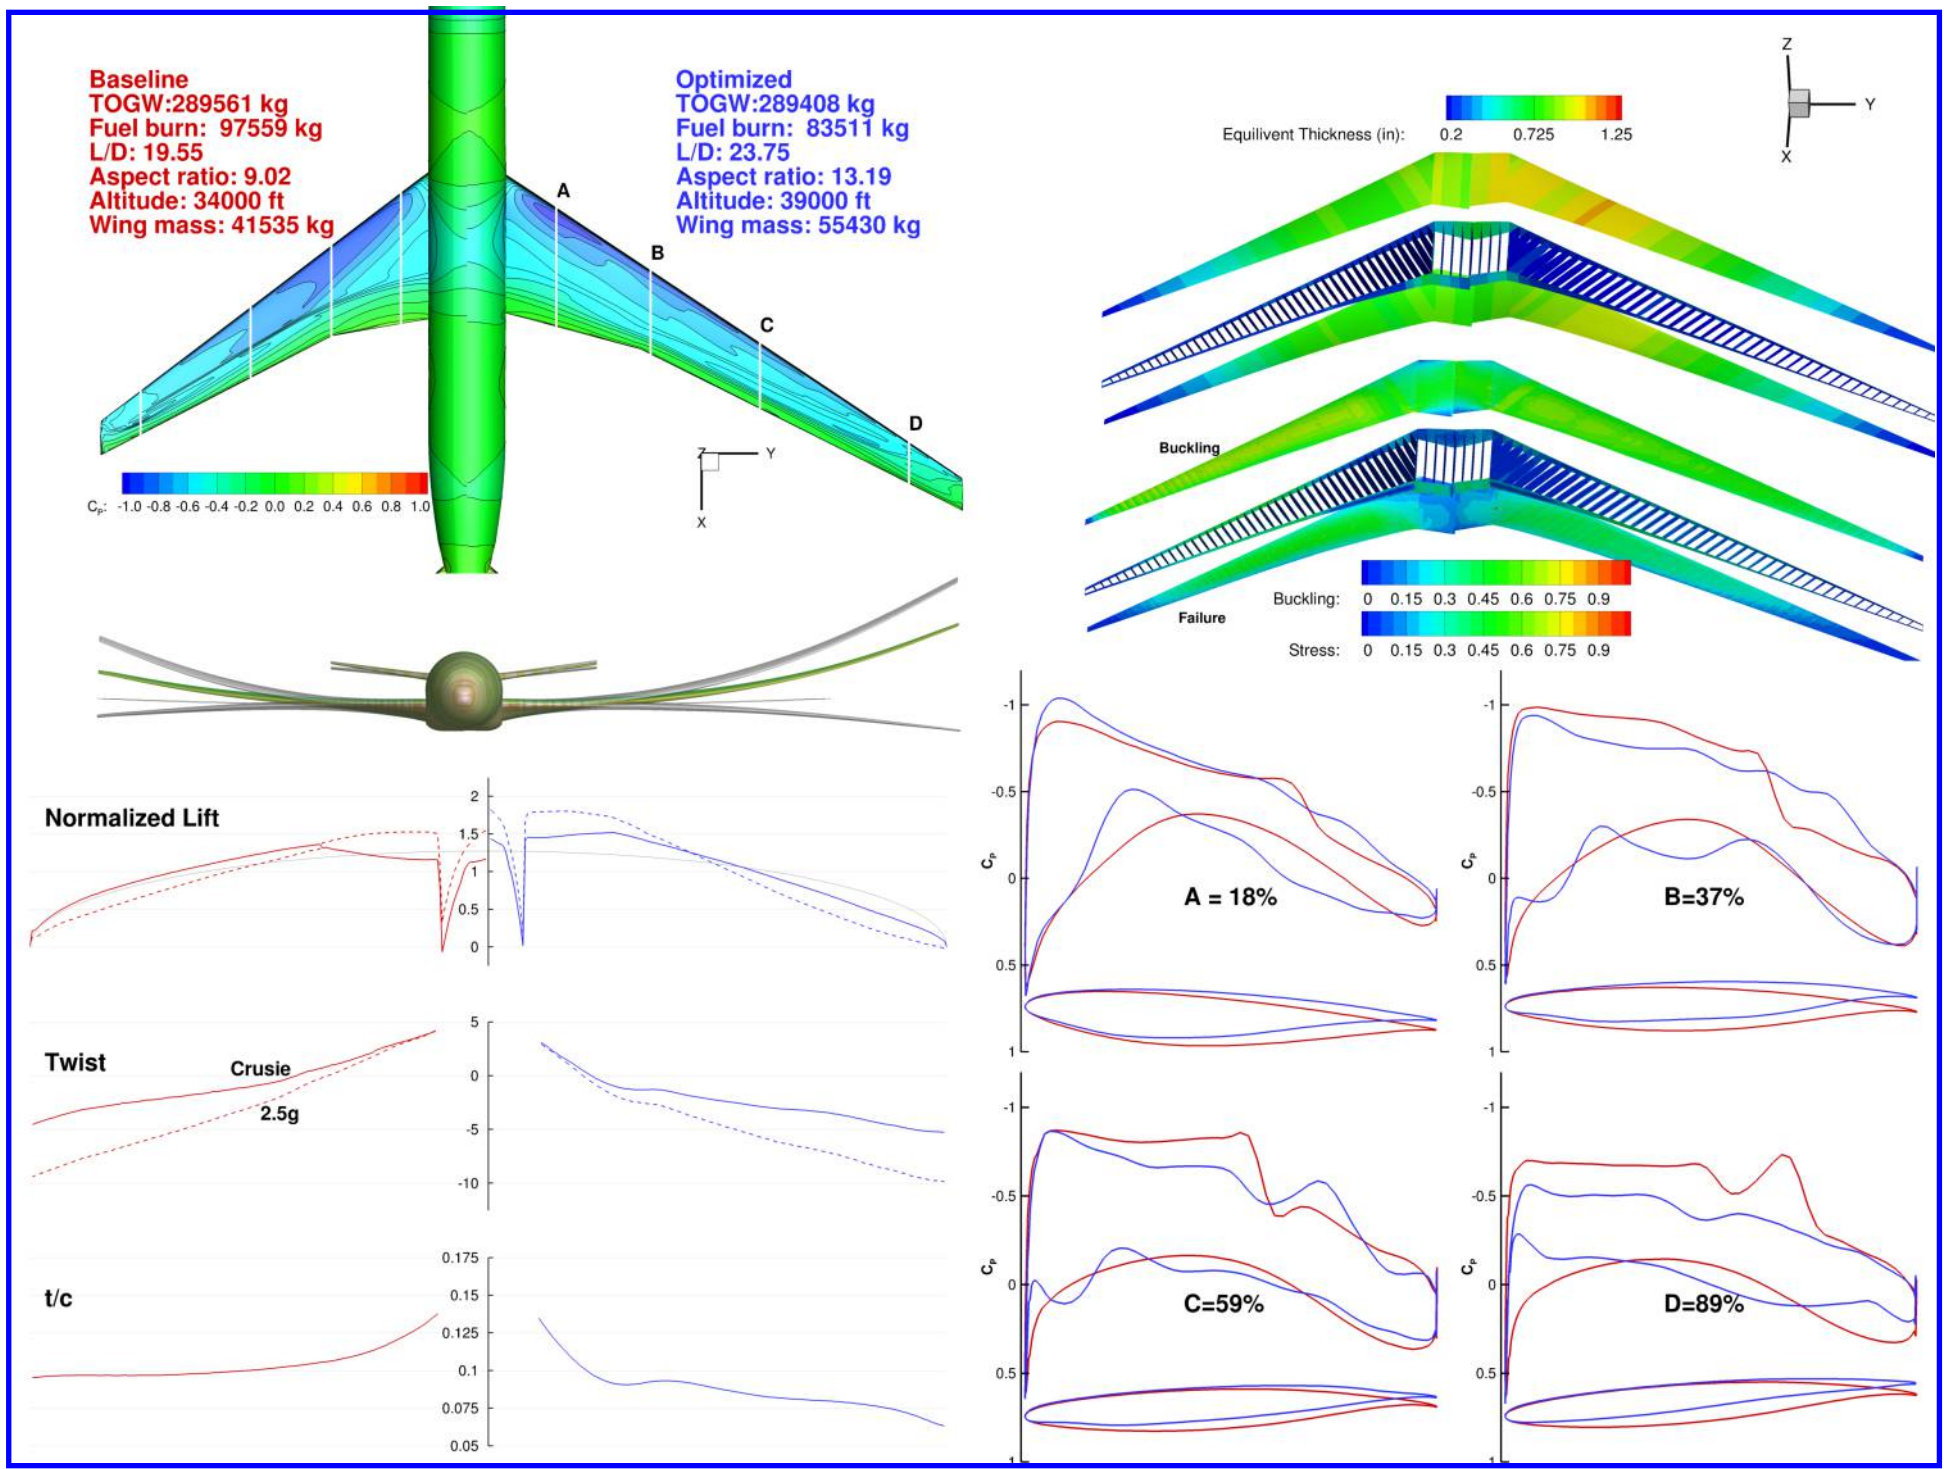
Figure 11. Overview of Case 1 optimization.
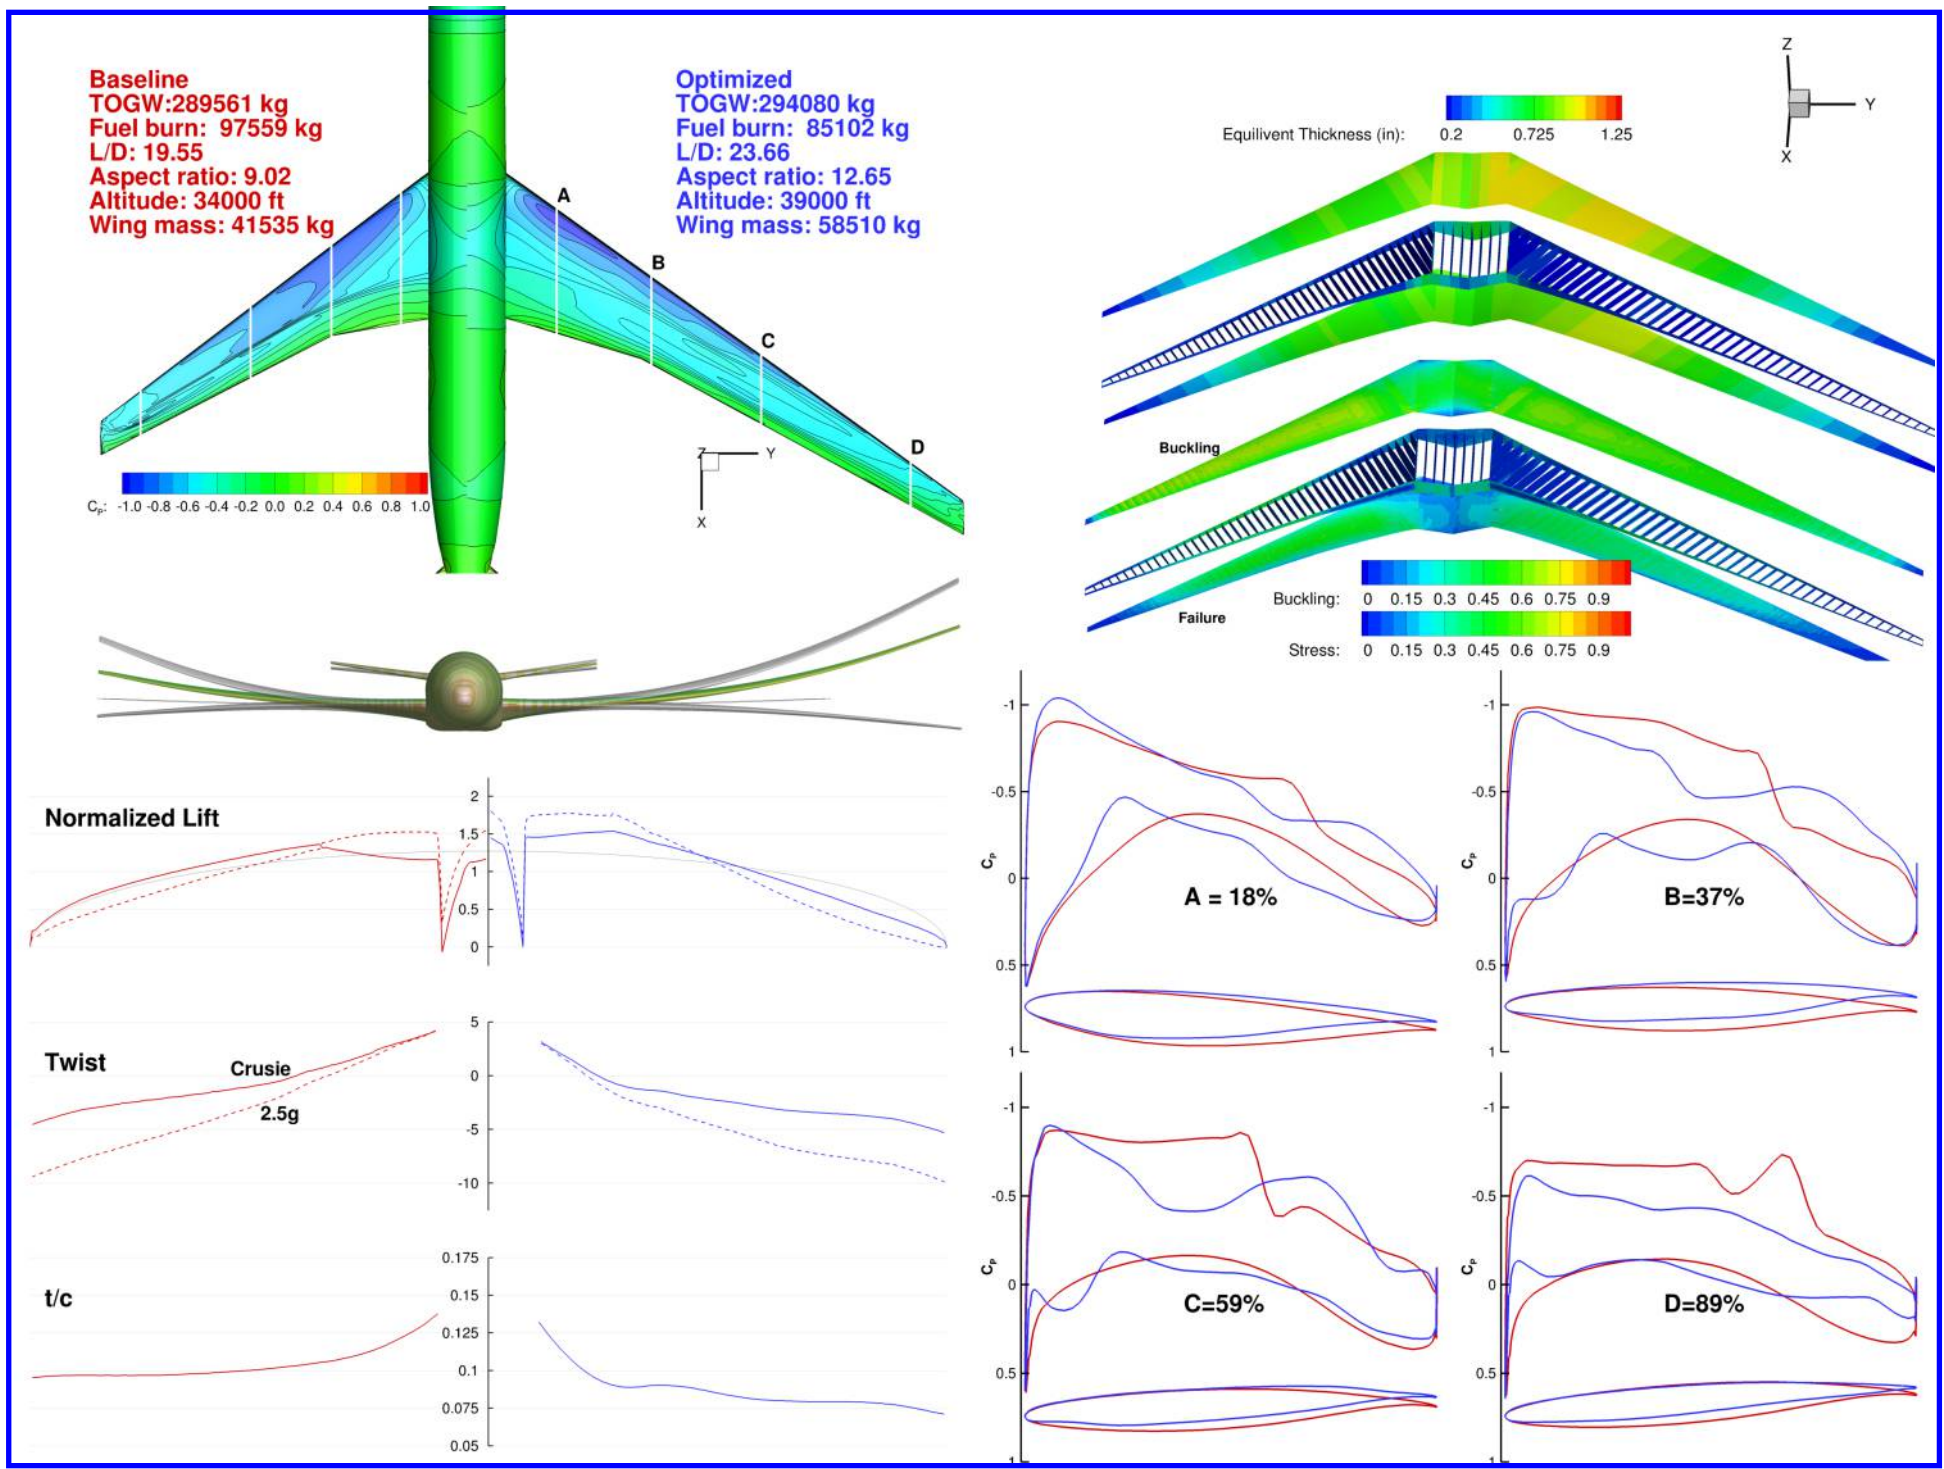
Figure 12. Overview of Case 2 optimization.

Figure 13. Overview of Case 3 optimization.

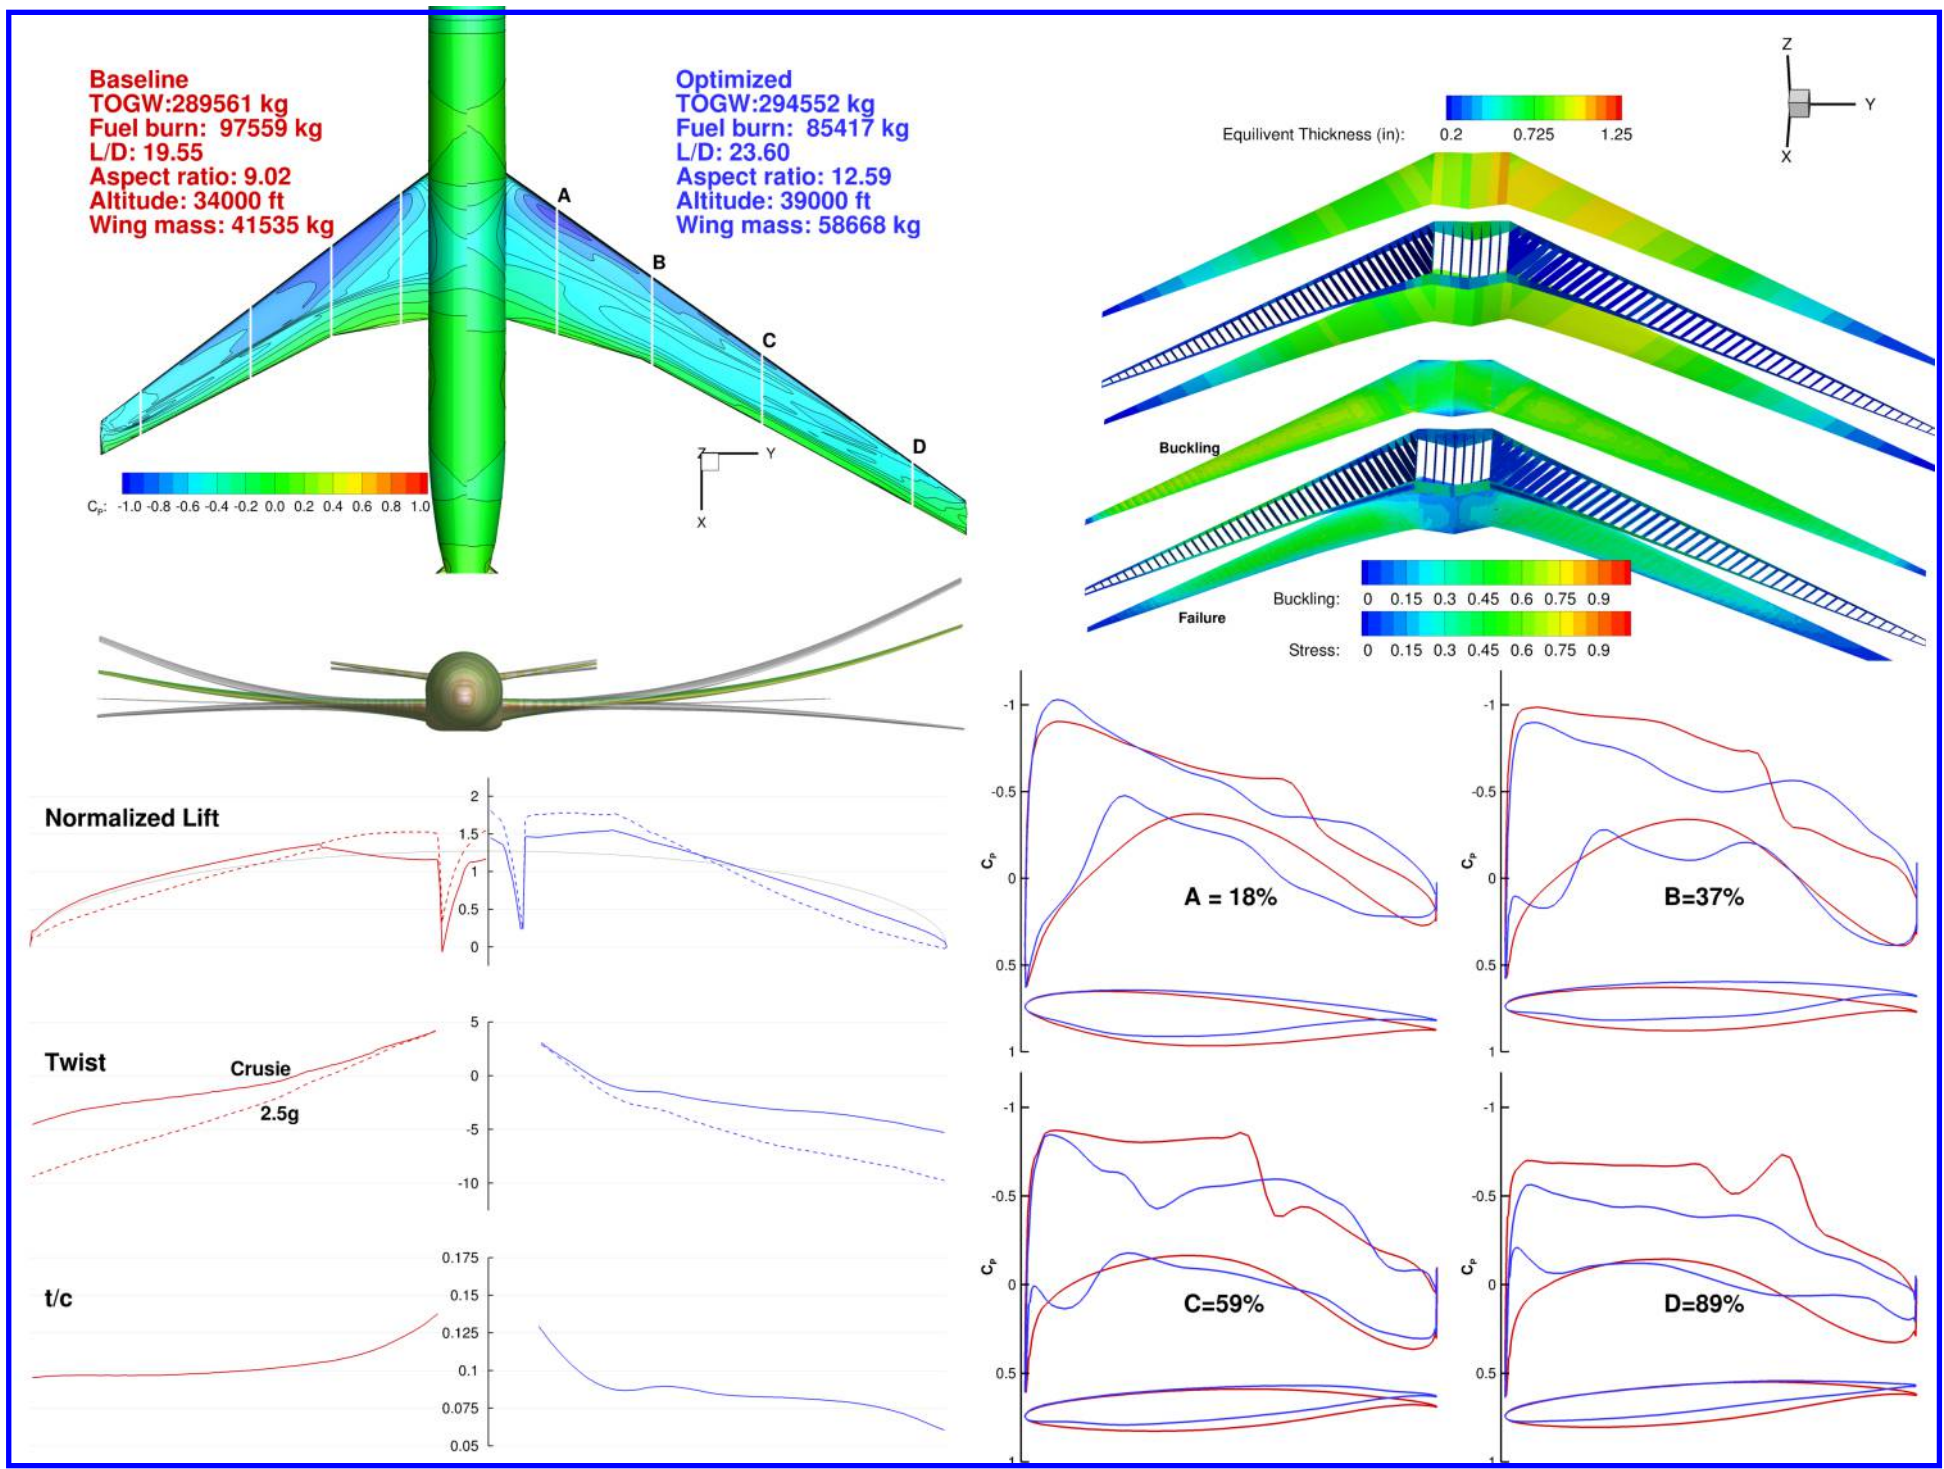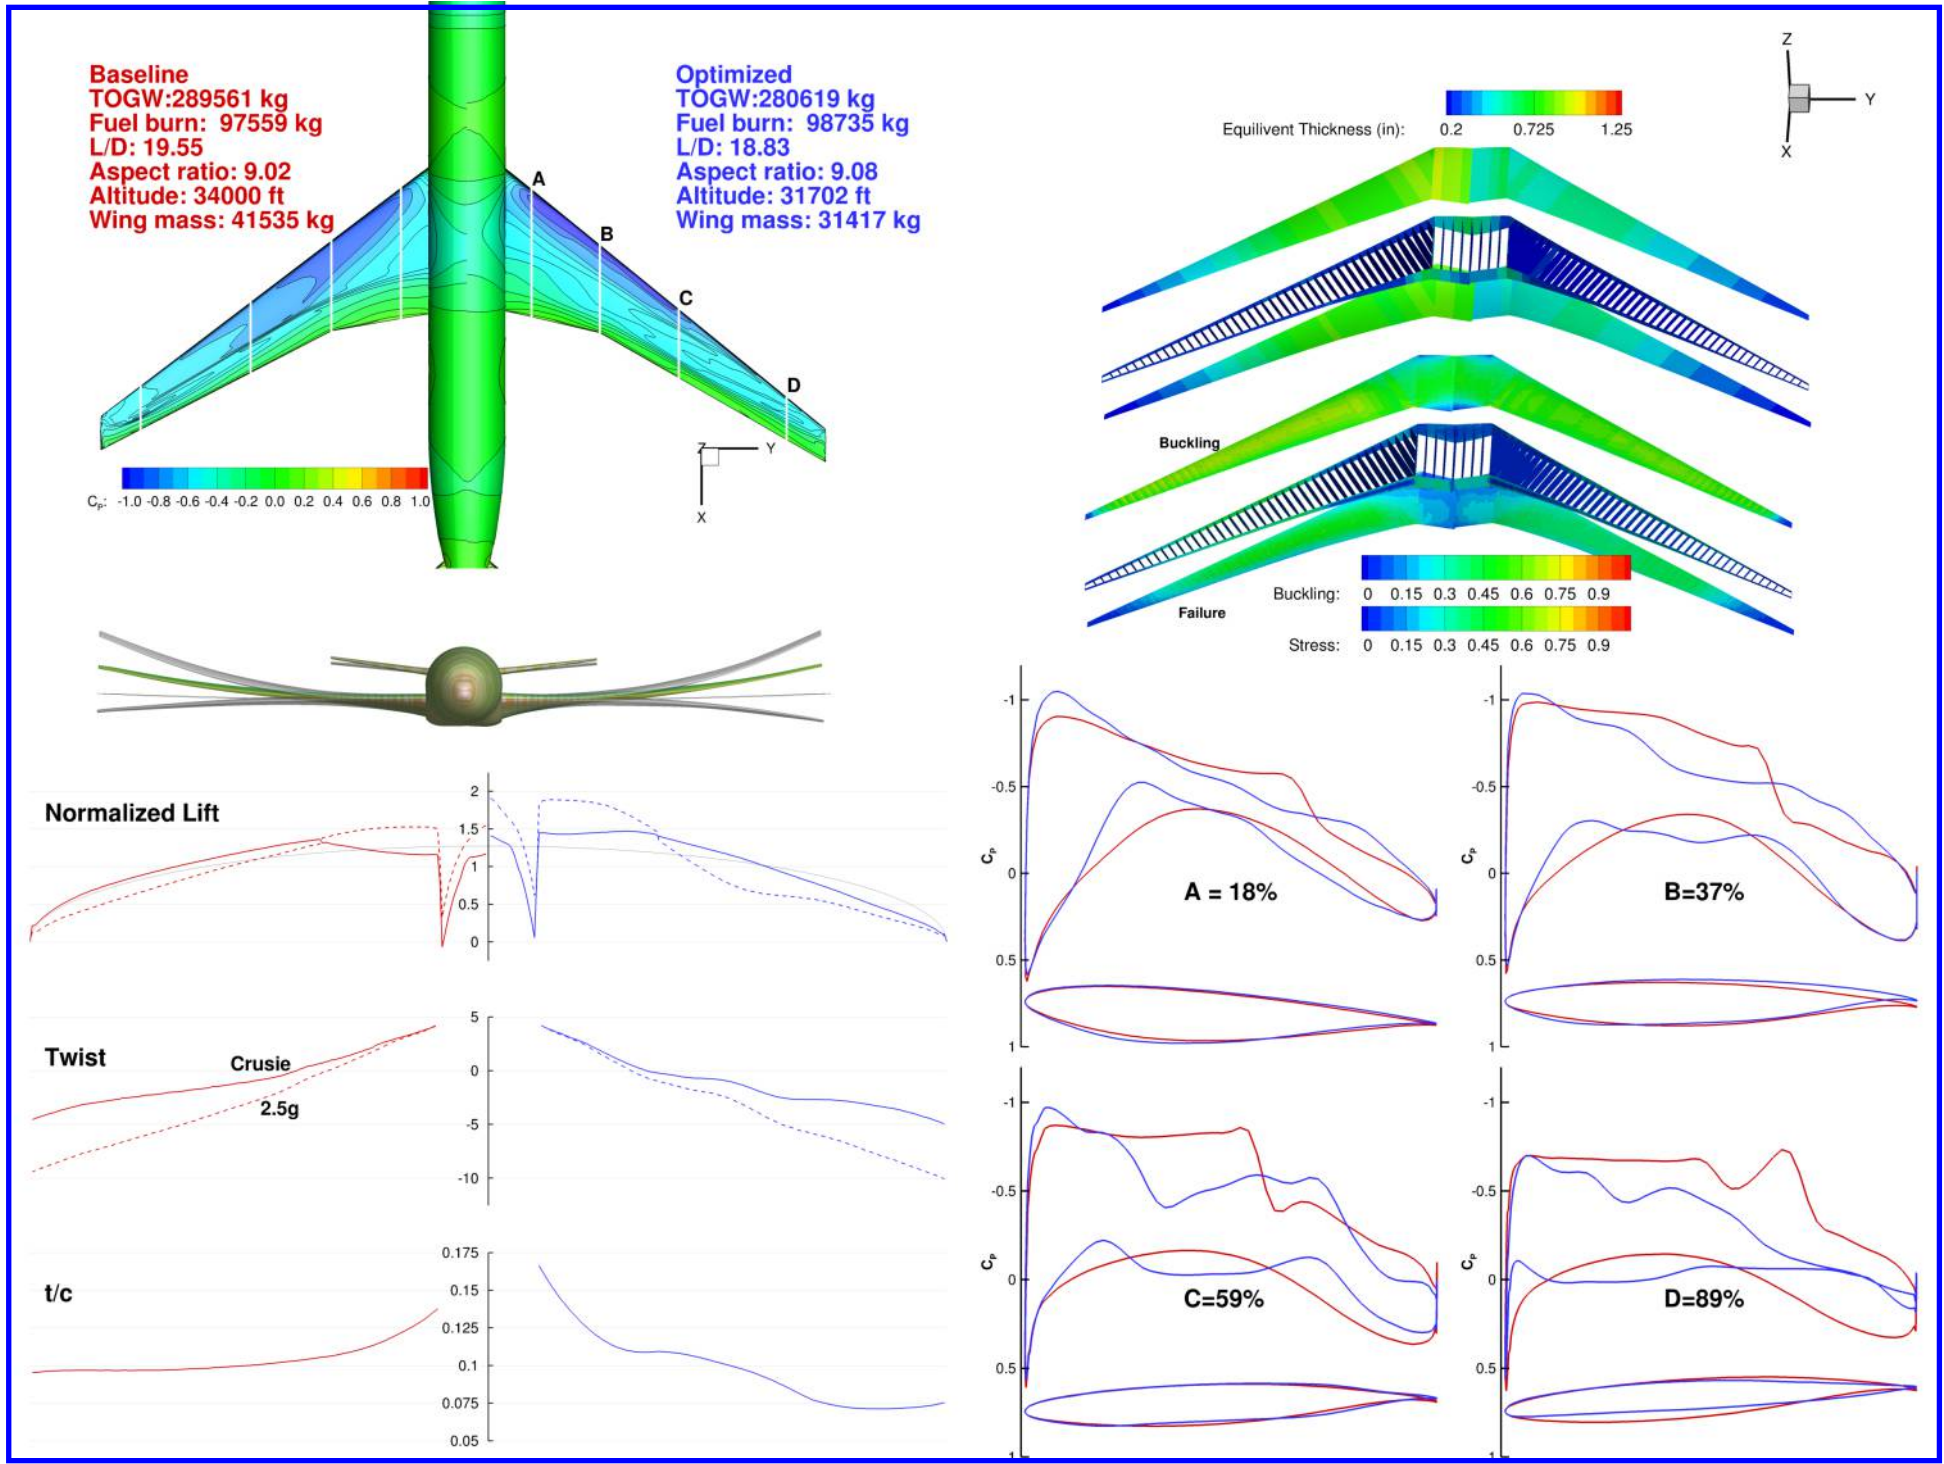**Baseline**
TOGW:289561 kg
Fuel burn: 97559 kg
L/D: 19.55
Aspect ratio: 9.02
Altitude: 34000 ft
Wing mass: 41535 kg

Optimized
TOGW:294552 kg
Fuel burn: 85417 kg
L/D: 23.60
Aspect ratio: 12.59
Altitude: 39000 ft
Wing mass: 58668 kg

A

B

C

D

C_p: -1.0 -0.8 -0.6 -0.4 -0.2 0.0 0.2 0.4 0.6 0.8 1.0

Equilvent Thickness (in): 0.2   0.725   1.25

Buckling

Failure

Buckling: 0 0.15 0.3 0.45 0.6 0.75 0.9

Stress: 0 0.15 0.3 0.45 0.6 0.75 0.9

Normalized Lift

Twist

Crusie

2.5g

t/c

C_p

A = 18%

C_p

B=37%

C_p

C=59%

C_p

D=89%

Figure 14. Overview of Case 4 optimization.

Figure 15. Overview of Case 4b optimization.

The Case 1 result sheds some light on the large changes observed in the structural mass. Despite the very large span increase, from 58.6 m to over 80 m, and a slight reduction in wing chord, the average t/c of the wing has been lowered considerably leading to an even larger wing box mass. As noted in previous papers [1, 14] the fuel burn objective places much more emphasis on lowering the wing drag. An objective sensitivity analysis for this configuration and mission shows that for the fuel burn objective, reducing the drag by 1 count is equivalent lowering the structural weight by approximately 850 kg. When this is considered, it is easy to see why the optimizer has increased the area (with a direct structural weight penalty) allowing the aircraft to operate in the thinner air at higher altitudes. Due to the significantly increased L/D of the optimized configurations and the fixed engine size, the excess thrust constraints, which would eventually limit the cruising altitude were never active over the course of the optimization. The equivalent wingbox thickness distribution in the upper right of Figures 11 through 14 confirms the wing skins are much thicker, resulting in the increased wing mass. The remaining fuel burn optimizations (cases 2 through 4) show similar overall trends: Large increases in area, lower thickness to chord ratios and thicker wing skins. The TOGW minimization, case 4b, shows very different trends. The wingspan is shortened slightly compared to the original configuration and the wingbox is thickened over the mid-span section. Outboard however, the wing is thinner than the baseline.

To gain a better understanding of how each of the optimized designs perform over a wide variety of transonic on-design and buffet onset conditions, a series of performance contour plots are generated. These are similar to the contours produced for aerodynamic optimizations performed previously by the authors [16]. These contour plots can be considered a two dimensional analogue of a drag polar.

We now briefly describe the procedure used to generate the contour plots of ML/cD in Figure 17. Their generation is complicated somewhat by the fact that all points need to be trimmed at a specified CG position. The procedure we employ is the same as the method provided by the DPW4 workshop organizing committee for producing a single trimmed drag polar. For completeness we summarize the procedure here.

A series of α sweeps were generated at fixed Mach numbers ranging from 0.8 to 0.88 and three specified tail angles: $-2°$, $0°$ and $2°$. The angle of attack range was sufficient to cover the range $C_L = 0.4$ to $C_L = 0.5$.

The trimming procedure is depicted graphically in Figure 16. To determine a single trimmed $M - C_L^*$ location, we first interpolate C_M and C_D at the desired C_L^* (left plot). After repeating the procedure for each tail angle, we plot the tail angles (η) as a function of the interpolated moments (C_M^*) (center plot). The tail angle required from trim, η^*, is then easily determined. Finally, the trimmed drag coefficient, $C_{D_{trim}}$, is determined by interpolating C_D^*, with the tail angle required for trim (Right plot).
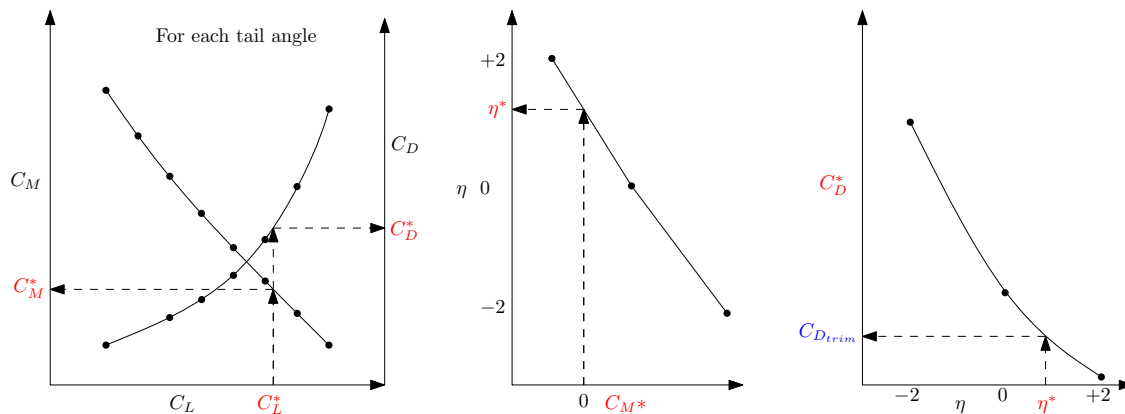

Figure 16. Procedure for determining drag at trim state.

Figure 17 also shows the buffet onset lines (in red) determined by the separation sensor method described in Section III. We also plot a 1.3 g buffet margin line (in blue). All normal flight operations must occur below this line. In addition, the position of the ML/cD maximum is denoted by a black square. A black contour at a value of 99% ML/cD$_{max}$ is also plotted. The motivation for plotting these 99% contours is that airliners typically fly between the Mach number yielding maximum range, (the maximum ML/cD value) and a higher Mach number that yields a 1% fuel-burn drag penalty but a decrease in the flight time. In addition, the area enclosed by the 99% contour is a measure of the on-design robustness: Any $M–C_L$ combination within the contour incurs no more than a 1% fuel-burn penalty. The operating conditions for the optimization are plotted with diamonds: The black diamond is the nominal operating condition, the blue diamonds are the additional normal operating conditions and the red diamonds are buffet onset conditions.
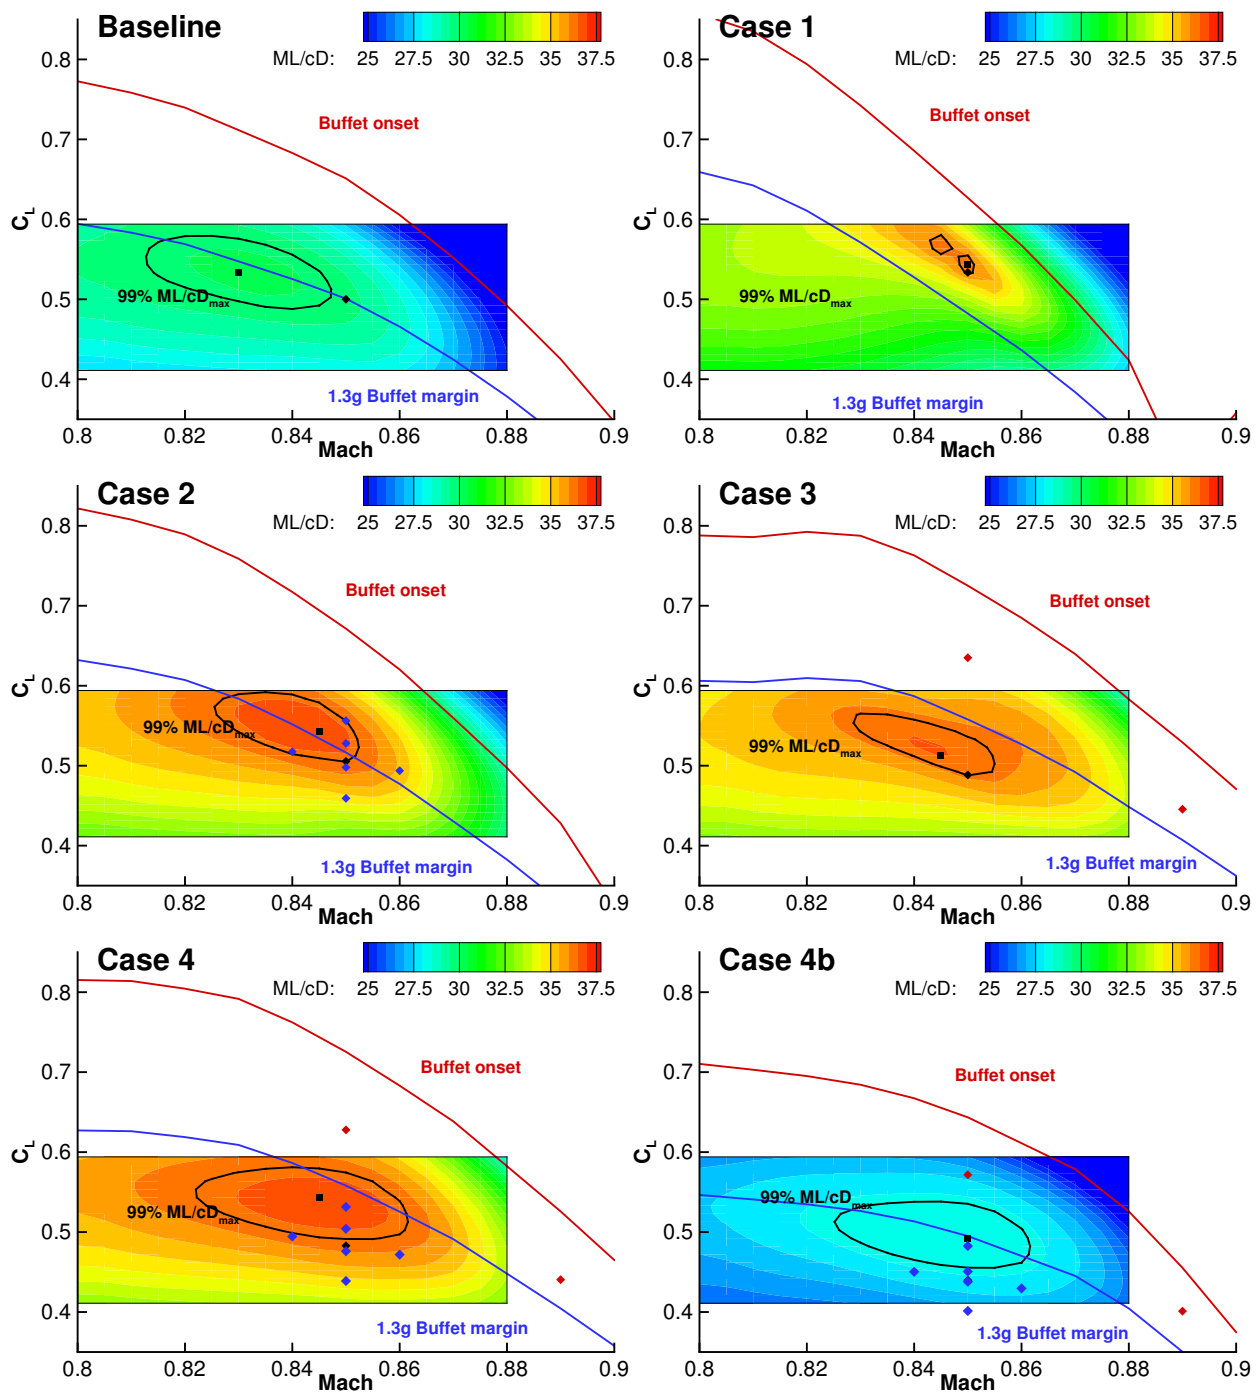
American Institute of Aeronautics and Astronautics

Figure 17. Including near buffet operating conditions results in more robust designs with a larger buffet margin.

Up until now, it appears that all the fuel burn optimizations produced fairly similar designs. A cursory analysis of the ML/cD contours in Figure 17 shows this is not the case. Case 1 is of particular interest. The 99% contour is both exceedingly small and multimodal. In addition, the entire high-performance region is completely above the 1.3 g buffet margin, indicating that the aircraft would not be allowed to operate in this region at all. It is clear this single point aerostructural optimization is insufficiently constrained. The on-design multipoint optimization of case 2 improves the situation considerably. For this case, the buffet boundary of the optimization design is similar to that of the initial design, and has even improved it somewhat over the lower Mach numbers. However, some of the assumed

American Institute of Aeronautics and Astronautics

operating conditions are over the 1.3 g margin curve and the point with maximum ML/cD is just outside the margin. This is obviously undesirable: we would very much like to operate the aircraft where it is most efficient. Moving on to case 3, the addition of the off-design buffet conditions has a significant impact on the design. Of particular interest is the vastly improved performance near the on-design condition, even though the two additional operating conditions are quite far away. For this case, the entire 99% ML/cD contour is within the 1.3 g boundary, however, the size of the contour is smaller than case 2, indicating less robust off-design performance near the operating condition. Case 4 which is a merging of cases 2 and 3, remedies both of the shortcomings of the previous two optimizations. The improved buffet boundary of case 3 is retained, while the robustness near the operating condition is further improved. For this design, the area enclosed inside the 99% contour is larger than for the baseline design. Finally, for case 4b, despite lowering of the on-design C_L the buffet conditions have ensured sufficient buffet margin, and the point of maximum ML/cD line is just below the 1.3 g margin curve. This is very similar to what is seen on the initial configuration.

Figure 18 gathers all the buffet boundaries, 99% $ML/cD_{\max}$ contours and maximum points onto the same plot. This allows for easy comparisons between the various designs. It is immediately clear that the inclusion of the near buffet conditions have a very positive impact on the buffet boundaries in Cases 3 and 4. These boundaries are everywhere higher than the initial design over this Mach number range. Case 2, even without the buffet conditions managed to maintain a boundary no worse than the initial with a slight improvement over the lower range of Mach numbers. Also, in addition to improving the selected objective, all optimizations increase the Mach number for $ML/cD_{\max}$ to between 0.845 and 0.85. The upper edge of the 99% contour for Cases 4 and 4b indicate efficient operation is possible at a Mach number as high as 0.86.
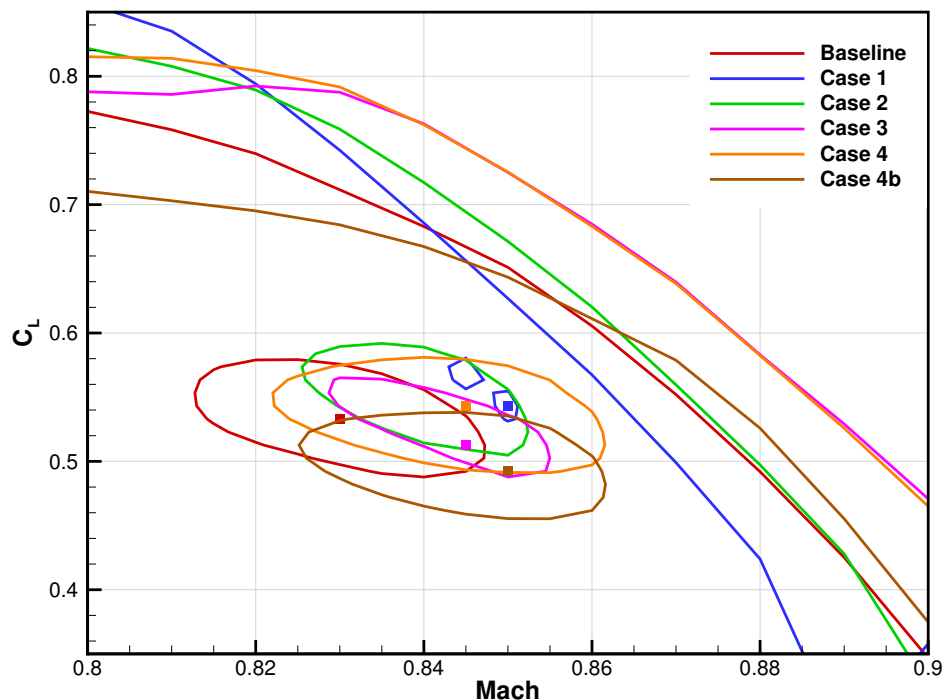

Figure 18. Comparison of buffet boundary and 99% ML/cD contours for each design.

Despite a fairly complete aerostructural optimization formulation, a thorough analysis of the optimized designs revealed several potential issues that may limit the practicality of the best fuel-burn optimized design, Case 4. Figure 19 shows the untrimmed $C_L - C_M$ curve for the baseline and optimized designs at the design Mach number of 0.85.

The slope of the curve gives an indication of the longitudinal static stability of the full configuration. The break in the pitching moment curve for the initial configuration at $C_L \sim 0.62$ is an indication of the appearance of separated flow and shows unstable post break-behavior. Note that this break in the pitching moment curve may also be used as indication of buffet onset. The case 4 design however, has a much different characteristic. Firstly, the overall slope in the "linear" region is substantially lower. The main reason for this is due to the large wing area increase without a corresponding increase in the tail area. Essentially, the pitching moment characteristics of the wing are overpowering
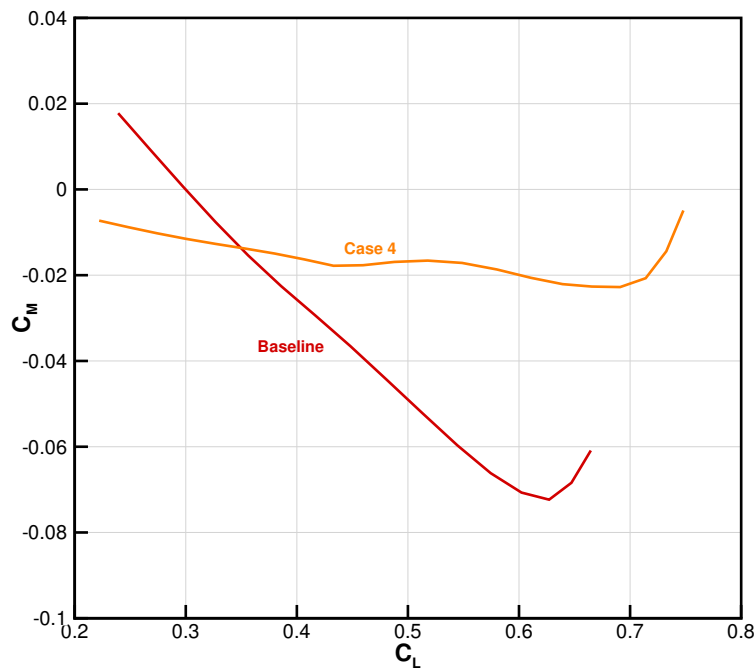
American Institute of Aeronautics and Astronautics

Figure 19. The optimized design has much lower static margin than the baseline design at Mach=0.85

the contribution from the tail. Secondly, there is a slightly unstable break near the on-design operating lift coefficients in the range of 0.4 to 0.55. While it might be possible for modern flight computers to handle such behavior, it is still undesirable. With a larger horizontal stabilizer, it is expected a more predictable behavior can be restored. We expect a larger horizontal stabilizer to only have small impact on the overall performance however. The induced drag penalty associated with trimming the larger wing is already accounted for with the current horizontal stabilizer. The additional drag associated with a larger stabilizer would come primarily from the increased skin friction drag due to the larger surface area only. This implies that the horizontal stabilizer design should be included in the aerostructural optimization. There is no fundamental limitation for doing so, except that a variety of other critical flight conditions that are affected by the stabilizer design must now be included in the optimization. This begins a snow-ball effect of ever increasing optimization complexity.

An alternative approach to restricting the optimization to only wing modifications would be to introduce a static stability constraint as was done in previous work [1]. This would naturally limit the wing area increase for a fixed sized horizontal stabilizer and would likely also prevent the design altitude from hitting the artificially defined limit.

On the structures side, with a significant increase in aspect ratio from 9.0 on the baseline design to approximately 12.6 for case 4, it is very possible that the three static maneuver conditions used for structural sizing are inadequate. Of particular concern is unsteady dynamic aeroelastic phenomena such as flutter.

There are additional practical considerations that may have a significant impact on the optimization. It is typical for wing mounted landing gear to be positioned at approximately the 50% MAC location. Normally, this is accommodated by extending the root chord and tapering the additional chord span-wise, resulting in the Yehudi break. Despite the fact that all the fuel burn optimized wings have lower sweep than the baseline configuration, the much higher aspect ratio has moved the leading edge MAC position rearward resulting in the situation shown in Figure 20. The addition of a landing gear post constraint along with separate chord variables for the root, Yehudi and tip chords, should result in a planform that satisfies this constraint.

VI.   Conclusions

The CRM has become a widely studied geometry for RANS based CFD analysis and aerodynamic design, and we are now starting to perform aerostructural studies the uCRM configuration. Using the uCRM geometry as a starting point, we performed high-fidelity aerostructural optimizations that are designed to investigate the engineering challenge of redesigning a wing for an existing aircraft while simultaneously improving the buffet onset margin.
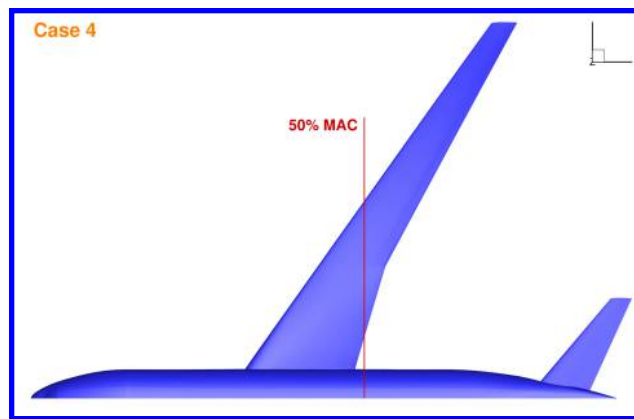
American Institute of Aeronautics and Astronautics

Figure 20. Wing mounted landing gear is not possible with this configuration.

A series of four increasing complexity fuel-burn minimization optimizations demonstrated the need for both multiple on-design operating conditions as well near-buffet conditions. The single point optimization, Case 1, reduced the fuel burn by 15.4%, but showed poor off-design behavior and an unacceptable buffet margin. The multipoint optimization with multiple near-design conditions achieved a 14.4% fuel burn reduction and vastly improved the off-design performance near the nominal operating condition. This set of operating conditions was sufficient to ensure the buffet boundary was no worse than the baseline, but the maximum performance condition was outside the 1.3 g buffet boundary. The addition of a 1.3 g and M_{MO} buffet conditions to the single point optimization significantly improved the both the buffet margin as well as the on-design robustness. This optimization reduced the fuel burn by 12.8%. The combination multiple on-design operating conditions with off design buffet-conditions (Case 4) yielded the lowest fuel burn reduction at 12.4%, but the highest robustness on-design and off-design performance. Significantly, the majority of the 99% ML/cD contour is available for operation due to the improved buffet boundary. Finally, Case 4b, using the TOGW objective instead of fuel burn, demonstrated the major impact of the choice of objective function on the overall design.

VII.   Acknowledgments

References

[1] Kenway, G. K. W. and Martins, J. R. R. A., "Multipoint High-Fidelity Aerostructural Optimization of a Transport Aircraft Configuration," *Journal of Aircraft*, Vol. 51, No. 1, January 2014, pp. 144–160. doi:10.2514/1.C032150.

[2] Kenway, G. K. W., Kennedy, G. J., and Martins, J. R. R. A., "Scalable Parallel Approach for High-Fidelity Steady-State Aeroelastic Analysis and Derivative Computations," *AIAA Journal*, Vol. 52, No. 5, May 2014, pp. 935–951. doi:10.2514/1.J052255.

[3] van der Weide, E., Kalitzin, G., Schluter, J., and Alonso, J. J., "Unsteady Turbomachinery Computations Using Massively Parallel Platforms," *Proceedings of the 44th AIAA Aerospace Sciences Meeting and Exhibit*, Reno, NV, 2006, AIAA 2006-0421.

[4] Mader, C. A., Martins, J. R. R. A., Alonso, J. J., and van der Weide, E., "ADjoint: An Approach for the Rapid Development of Discrete Adjoint Solvers," *AIAA Journal*, Vol. 46, No. 4, April 2008, pp. 863–873. doi:10.2514/1.29123.

[5] Spalart, P. and Allmaras, S., "A One-Equation Turbulence Model for Aerodynamic Flows," *30th Aerospace Sciences Meeting and Exhibit*, 1992. doi:10.2514/6.1992-439.

[6] Lyu, Z., Kenway, G. K., Paige, C., and Martins, J. R. R. A., "Automatic Differentiation Adjoint of the Reynolds-Averaged Navier–Stokes Equations with a Turbulence Model," *21st AIAA Computational Fluid Dynamics Conference*, San Diego, CA, Jul 2013. doi:10.2514/6.2013-2581.

[7] Luke, E., Collins, E., and Blades, E., "A Fast Mesh Deformation Method Using Explicit Interpolation," *J. Comput. Phys.*, Vol. 231, No. 2, Jan. 2012, pp. 586–601. doi:10.1016/j.jcp.2011.09.021.

[8] 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, Sept. 2010, AIAA 2010–9308.

[9] Kennedy, G. J. and Martins, J. R. R. A., "A parallel aerostructural optimization framework for aircraft design studies," *Structural and Multidisciplinary Optimization*, Vol. 50, No. 6, December 2014, pp. 1079–1101. doi:10.1007/s00158-014-1108-9.

[10] Brown, S. A., "Displacement Extrapolation for CFD+CSM Aeroelastic Analysis," *Proceedings of the 35th AIAA Aerospace Sciences Meeting*, Reno, NV, 1997, AIAA 1997-1090.

[11] Irons, B. M. and Tuck, R. C., "A Version of the Aitken Accelerator for Computer Iteration," *International Journal for Numerical Methods in Engineering*, Vol. 1, No. 3, 1969, pp. 275–277.

[12] Gill, P. E., Murray, W., and Saunders, M. A., "SNOPT: An SQP algorithm for large-scale constrained optimization," *SIAM Journal of Optimization*, Vol. 12, No. 4, 2002, pp. 979–1006. doi:10.1137/S1052623499350013.

[13] Perez, R. E., Jansen, P. W., and Martins, J. R. R. A., "pyOpt: a Python-Based Object-Oriented Framework for Nonlinear Constrained Optimization," *Structural and Multidisciplinary Optimization*, Vol. 45, No. 1, January 2012, pp. 101–118. doi:10.1007/s00158-011-0666-3.

[14] Kenway, G. K. W., Kennedy, G. J., and Martins, J. R. R. A., "Aerostructural optimization of the Common Research Model configuration," *15th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*, Atlanta, GA, June 2014, AIAA 2014-3274.

[15] Lyu, Z. and Martins, J. R. R. A., "Aerodynamic Shape Optimization of an Adaptive Morphing Trailing Edge Wing," *Journal of Aircraft*, 2015, (Accepted).

[16] Lyu, Z., Kenway, G. K., and Martins, J. R. R. A., "Aerodynamic Shape Optimization Studies on the Common Research Model Wing Benchmark," *AIAA Journal*, Vol. 53, No. 4, April 2015, pp. 968–985. doi:10.2514/1.J053318.

[17] Mader, C. A. and Martins, J. R. R. A., "Stability-Constrained Aerodynamic Shape Optimization of Flying Wings," *Journal of Aircraft*, Vol. 50, No. 5, September 2013, pp. 1431–1449. doi:10.2514/1.C031956.

[18] Obert, E., *Aerodynamic Design of Transport Aircraft*, IOS Press BV, 2009.

[19] Lovely, D. and Haimes, R., "Shock detection from computational fluid dynamics results," *Proceedings of the 14th Computational Fluid Dynamics Conference*, Norfolk, VA, January 1999. doi:doi:10.2514/6.1999-3285.

[20] Vassberg, J. C., DeHaan, M. A., Rivers, S. M., and Wahls, R. A., "Development of a Common Research Model for Applied CFD Validation Studies," 2008, AIAA 2008-6919.

[21] Lyu, Z., Kenway, G. K. W., and Martins, J. R. R. A., "RANS-based Aerodynamic Shape Optimization Investigations of the Common Research Model Wing," *Proceedings of the AIAA Science and Technology Forum and Exposition (SciTech)*, National Harbor, MD, January 2014. doi:10.2514/6.2014-0567, AIAA 2014-0567.

[22] Vassberg, J. and Jameson, A., "Influence of Shape Parameterization on Aerodynamic Shape Optimization," Tech. rep., Von Karman Institute, Brussels, Belgium, April 2014.

[23] Telidetzki, K., Osusky, L., and Zingg, D. W., "Application of Jetstream to a Suite of Aerodynamic Shape Optimization Problems," *52nd Aerospace Sciences Meeting*, Feb 2014. doi:10.2514/6.2014-0571.

[24] Carrier, G., Destarac, D., Dumont, A., Meheut, M., Din, I. S. E., Peter, J., Khelil, S. B., Brezillon, J., and Pestana, M., "Gradient-Based Aerodynamic Optimization with the elsA Software," *52nd Aerospace Sciences Meeting*, Feb 2014. doi:10.2514/6.2014-0568.

[25] Anonymous, "777-200/300 Airplane Characteristics for Airport Planning," Tech. Rep. D6-58329, Boeing Commercial Airplanes, July 1998.

American Institute of Aeronautics and Astronautics

This article has been cited by:

1. R. L. T. Bevan, D. J. Poole, C. B. Allen, T. C. S. Rendall. 2017. Adaptive Surrogate-Based Optimization of Vortex Generators for Tiltrotor Geometry. *Journal of Aircraft* **54**:3, 1011-1024. [Abstract] [Full Text] [PDF] [PDF Plus]
2. Gaetan K. W. Kenway, Joaquim R. R. A. Martins. 2016. Multipoint Aerodynamic Shape Optimization Investigations of the Common Research Model Wing. *AIAA Journal* **54**:1, 113-128. [Abstract] [Full Text] [PDF] [PDF Plus] [Supplemental Material]