Buffet Onset Constraint Formulation for Aerodynamic Shape Optimization

Gaetan K. W. Kenway[1]
Joaquim R. R. A. Martins[2]
University of Michigan, Department of Aerospace Engineering, Ann Arbor, MI

Abstract High-fidelity computational modeling and optimization of aircraft configurations have the potential to enable engineers to create more efficient designs that require fewer unforeseen modifications late in the design process. Although aerodynamic shape optimization has the potential to produce high-performance transonic wing designs, these designs remain susceptible to buffet. To address this issue, we present a separation-based constraint formulation that constrains buffet onset in an aerodynamic shape optimization. The separation metric developed in this work is verified against a common buffet prediction method and validated against experimental wind tunnel data. A series of optimizations based on the AIAA Aerodynamic Design Optimization Discussion Group wing-body-tail case are presented to show that buffet-onset constraints are required and to demonstrate the effectiveness of the proposed approach. Although both single-point and multipoint optimizations without separation constraints are vulnerable to buffeting, the optimizations using the proposed approach move the buffet boundary to make the designs feasible.

1 Introduction

Numerical optimization is a powerful tool that can complement more traditional design methodologies. Design optimization based on high-fidelity, physics-based simulations, such as computational fluid dynamics (CFD) and computational structural mechanics, is especially promising [1, 2]. By capturing the relevant physics of the underlying system, performance improvements predicted by numerical simulations are more likely to be realized in the real world. Effective optimization algorithms, however, invariably exploit limitations in the numerical models or in incomplete formulations of the optimization problems by violating important design constraints that are not included.

CFD-based aerodynamic shape optimization dates back to Hicks *et al.* [3], who first tackled airfoil design optimization problems, and has steadily evolved over the last few decades. One of the major advancements in this field was the development of adjoint methods [4, 5, 6], which in conjunction with gradient-based optimization has enabled optimization with respect to large numbers of shape parameters, enabling the aerodynamic shape optimization of wings [7, 8, 9, 10, 11, 12, 13, 14] and full configurations [15, 16, 17]. Recently, a series of benchmarks that include airfoil, wing, and wing-body cases was developed by the AIAA Aircraft Design Optimization Discussion Group (ADODG). These cases allow researchers to compare the results of different design methods and to evaluate their relative strengths and weaknesses [12, 18, 19, 20, 21, 13]. The adjoint technique used in aerodynamic shape optimization has also been extended to simulations that couple aerodynamics and structures, enabling the simultaneous design optimization of outer

mold line shape and structural sizing, while accounting for wing flexibility and aerostructural design trade-offs [2, 1, 22, 23]. Considering multiple disciplines in design optimization is the basic principle of multidisciplinary design optimization (MDO) [24], and the coupled adjoint technique has been extended to general MDO problems [25].

Buffet is a critical aspect of transonic wing design and has yet to be explicitly considered in aerodynamic shape optimization. Buffet may be broadly defined as a high-frequency aerodynamic instability caused by flow separation. This instability is undesirable because the resulting unsteady aerodynamic loads compromise the ability to comfortably control the aircraft and the aerodynamic performance.

In the transonic regime, as the Mach number or lift coefficient increases, shocks on the wing gradually increase in strength. The interaction of the shock-induced separation at the foot of the shock with the oscillation of the shock causes *transonic buffet*, which limits the maximum aircraft lift coefficient and Mach number. Because jet transport aircraft are most efficient when a combination of cruise Mach number and lift-to-drag ratio is maximized, buffet is often an active design constraint. The maximum lift coefficient at a given Mach number decreases with increasing Mach number, effectively limiting the aircraft altitude, i.e., the *aerodynamic ceiling*. This ceiling is also known as the "coffin corner," because, although the Mach number may be decreased to increase the lift coefficient and therefore achieve a higher altitude, the stall speed also increases due to the decrease in density, which reduced the difference between the stall speed and maximum speed. Ultimately, the stall and buffet boundaries intersect at a sufficiently high altitude, making the aircraft impossible to fly—hence the "coffin corner."

Given that buffet crucially affects transport aircraft performance, a need exists for an effective way to formulate buffet as a design constraint. Although buffet has been considered in a few design optimization studies, it has yet to be considered as a constraint in CFD-based design optimization. Wakayama *et al.* [26] performed MDO of a wing using a low-fidelity method calibrated against CFD to estimate buffet onset based on Mach number, local wing sweep, thickness-to-chord ratio, and lift coefficient. The method was calibrated against flight test and CFD calculations, and it was also used in the MDO of a blended-wing body configuration [27]. More recently, Bérard and Isikveren [28] have developed another inexpensive approach to enforce a buffet onset constraint in aircraft conceptual design optimization. Buffet has been mentioned in the context of CFD-based aerodynamic shape optimization as a requirement to verify after optimization [29, 30], and it has also been implicitly considered in airfoil optimization by the addition of the drag at off-design conditions to the objective function [31]. Thus, a need exists to develop a CFD-based method to explicitly enforce a buffet onset constraint.

By using unsteady CFD, a number of researchers have made strides towards modeling the physics of transonic buffet of airfoils [32, 33, 34, 35, 36, 37, 38]. However, unsteady CFD is currently too computationally intensive to serve as a constraint in a design optimization because it requires hundreds of objective and constraint function evaluations. To address this issue, Thomas and Dowell [39] used the frequency domain approach to model the unsteady aerodynamics; a technique that was previously used in design optimization involving unsteady phenomena [40, 41]. They also implemented a discrete adjoint to obtain gradients and demonstrated the use of this approach in the optimization of an airfoil. Although they did not implement buffet onset as a constraint, they minimized the peak of the unsteady loading for an NACA 0012 airfoil.

In the present work, we are not interested in modeling actual unsteady transonic flow shock buffeting. Instead, our goal is to predict the transonic buffet onset lift coefficient for fixed Mach and Reynolds numbers, so that we can add a design constraint that keeps the wing design within the buffet boundary at $1.3g$. If this constraint were implemented correctly, the optimization would minimize the drag at the design lift coefficients, subject to buffet constraints. It is particularly

important to implement the buffet requirement as a true optimization constraint, as opposed to adding buffet off-design points in the drag minimization because (1) the optimal solution might not actually satisfy the buffet requirement, and (2) the solution will be suboptimal with respect to the properly constrained formulation.

To quantify buffet onset in aerodynamic shape optimization, we develop a new prediction method that is based on the extent of separated flow present on the wing in a steady Reynolds-averaged Navier–Stokes (RANS) CFD computation. We then use this method constrain the extent of separated flow near the buffet onset boundary, which ensures that the optimized design has a sufficient buffet margin, while simultaneously improving the performance at the design operating conditions. The constraint function is smooth and its gradient with respect to the wing shape variables is computed by using a discrete adjoint method. We demonstrate the effectiveness of the proposed approach by applying it to the multipoint drag minimization of the wing-body-tail geometry defined by ADODG Case 5 [21].

This paper is organized as follows: We begin by outlining the key aspects of the computational methods used in this work, following which we describe the separation constraint formulation and how it is used to enforce the buffet onset constraint. We verify the proposed approach by comparing it with the results of an alternate numerical approach, and we also validate it by comparing it with the results of a wind tunnel experiment. Finally, we present a sequence of numerical aerodynamic optimization studies based on the ADODG wing-body-tail case to evaluate how the buffet-onset constraints affect in transonic wing aerodynamic shape optimization.

2 Computational Methods

In this work, the aerodynamic shape optimization is done with the MDO of Aircraft Configuration with High Fidelity (MACH) framework. This framework was developed for the aerostructural design optimization of aircraft configurations [2, 1], and it integrates modules for CFD, structural analysis, geometry parametrization, and numerical optimization. MACH has been used extensively for both aerodynamic shape [11, 12, 42, 17, 14] and aerostructural design optimization [43, 23, 44, 45] of aircraft and hydrofoils [46]. Herein, we use only the aerodynamic capabilities of MACH, which we describe in the remainder of this section.

1 Computational Fluid Dynamics Solver

The flow solver in MACH is ADflow, which solves the RANS equations in either steady, unsteady, or time spectral modes [47, 48]. ADflow applies the finite-volume method to structured, body fitted, multiblock grids. The discretization scheme uses central fluxes with artificial dissipation and the Spalart–Allmaras turbulence model [49]. A matrix dissipation scheme [50] is used herein except where explicitly noted. A fully coupled Newton–Krylov method is used to simultaneously solve the mean flow and turbulence equations. A discrete adjoint method is implemented by using a combination of reverse-mode automatic differentiation and analytic methods for the efficient computation of the gradients of functions of interest. Lyu *et al.* [51] describe the CFD adjoint implementation in more detail.

2 Geometric Parametrization

In this work, we use a free-form deformation (FFD) volume approach [52] that we implemented [53] and have used extensively in the past for aerodynamic [10, 54, 11, 12, 17, 55] and aerostructural optimization studies [2, 43, 1, 23]. The FFD approach may be visualized as embedding the spatial coordinates that define a geometry inside a flexible volume. The parametric locations corresponding to the baseline geometry are found by using a Newton search algorithm. Once the baseline geometry

is embedded, perturbations made to the FFD volume propagate within the embedded geometry by evaluating the nodes at their parametric locations.

3   Mesh Movement

The FFD approach used to parametrize the geometry applies deformations only to the surface mesh; that is, the part of the volume mesh that lies on the physical surface. A separate procedure is then required to propagate surface perturbations throughout the remainder of the volume mesh. The mesh movement algorithm used in this work is an efficient analytic inverse distance method similar to that described by Luke *et al.* [56]. Updating the mesh for a new configuration is fast, typically requiring less than 0.1% of the CFD solution time. Sensitivities required for the adjoint method are provided by a combination of reverse-mode automatic differentiation and analytic methods.

4   Optimization Algorithm

The high computational cost of RANS-based optimization demands an optimization algorithm that minimizes the number of function evaluation calls. We use SNOPT (sparse nonlinear optimizer) [57] with the Python interface pyOpt [58]. 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 already used the SNOPT algorithm to solve a wide variety of aerodynamic and aerostructural optimization problems [1, 23, 42, 12, 10].

3   Buffet-Onset Prediction

In the broadest sense, buffet is any form of vibration caused by unsteady forces generated by separated flow. There are three main types of flow separation: (1) separation at the foot of a shock wave, (2) leading-edge separation, and (3) trailing-edge separation. Transonic (or high-speed) buffet is caused by the first type: shock-induced separation. In transonic flow, at sufficiently high lift coefficients and Mach numbers, instabilities in the interaction between the shock and the separation bubble cause self-sustaining periodic oscillations in the shock position [38], which cause large fluctuations in pressure with a frequency on the order of 10 Hz [59]. Buffet is primarily an aerodynamic phenomenon because the frequencies of the shock-induced vibrations are at least one order of magnitude greater than the natural frequencies of the wing's primary elastic modes, so no aeroelastic computations are required to predict it. Buffet develops gradually with increasing lift coefficient or Mach number, and *buffet onset* refers to the conditions at which buffet first occurs. As the Mach number increases, the buffet-onset lift coefficient decreases, defining the *buffet boundary*.

Buffet is undesirable because it affects the ability to control the aircraft and passenger comfort, and, if severe enough, it may compromise the structural integrity of the aircraft. Therefore, Joint Aviation Requirements stipulate that commercial transport aircraft maintain at least a 30% margin from the cruise operating condition to buffet onset. This buffet margin provides a margin of maneuverability for the aircraft. This allows the aircraft to perform a $1.3g$ maneuver in cruise flight, which is equivalent to turning at a $40°$ bank angle. In addition to ensuring that maneuvers can be executed free of buffet, this margin also ensures that disturbances due to turbulence and upsets due to aircraft system failures can be handled safely. In this work, we seek a way to predict the buffet onset numerically, so that our aerodynamic shape optimization can stay within the boundary defined by the 30% margin.

Numerous researchers have modeled transonic buffet for airfoils with unsteady CFD using large eddy simulations [35], detached eddy simulations [32] and unsteady RANS [33, 34, 36, 37, 38].

4

Although such simulations have the merit of clarifying the physics, time-accurate CFD is currently too computationally costly to include in a numerical optimization process, because hundreds of such time-accurate simulations would be required to complete an optimization. On the low-fidelity side, buffet constraints have been implemented in conceptual aircraft design optimization [26, 27, 28]. However, these low-fidelity methods do not consider how the detailed airfoil shape design affects buffet onset, which is critical when performing CFD-based aerodynamic shape optimization.

Therefore, the goal of the present work is to develop a CFD-based method to explicitly enforce a buffet-onset constraint. To achieve this, we do not need to model the unsteady transonic flow and the physics beyond buffet onset. Instead, by using steady RANS data, we formulate a constraint function that predicts if a design is within the buffet boundary . Although the literature is split on whether or not steady RANS accurately models transonic buffet, Rumsey et al. [60] report steady RANS predictions that are consistent with flight data through the buffet onset regime and up to near the maximum lift coefficient. However, they stipulate that this agreement is not well understood and might be case-dependent. To clarify this question, we present below verification and validation results that confirm that steady RANS is well suited for our purposes.

To ensure that a gradient-based optimization algorithm can handle the buffet constraint, the constraint function should be continuous and change smoothly upon approaching the buffet boundary. Although the actual physical behavior is highly nonlinear, buffet onset is a gradual process, so developing such a function should be possible.

A number of techniques have been developed over the years to correlate data, typically from wind tunnel experiments, with buffet predictions from a flight testing program. These techniques include correlations with the RMS signals from wind tunnel model strain gauges, trailing-edge pressure divergence, axial force break, pitching moment break, and lift curve break [61, 62]. The last two methods may be employed in numerical predictions by using CFD to integrate force and moment values. One way to implement the lift curve-break method is the $\Delta\alpha = 0.1$ method [63]. Using this method, the linear portion of the lift curve is offset to the right by $0.1°$. The intersection of this line with the actual lift curve is used to estimate the buffet onset point, as illustrated in Fig. 1. We could use this method to develop a buffet-onset constraint function, but in using global aerodynamic coefficients, such as C_L and C_M, we would not make full use of the detailed flow solution provided by CFD. In addition, the "linear" portion of the lift curve slope is not exactly linear in transonic flow, and identifying the slope to be used introduces ambiguity. The use of this approach with CFD requires at least two additional flow solutions (one for the slope and one for the intersection). Finally, implementing a constraint based on this method would provide the optimizer with the opportunity to artificially affect the buffet onset by manipulating the lift curve at lower lift coefficients.

To develop a more direct way of constraining buffet onset, we focus on the physical mechanism of shock-induced flow separation, which is responsible for the loss of lift and the subsequent lowering of the lift curve slope. An example showing the typical progression of this type of separation with increasing angle of attack is shown in Fig. 2. To obtain the results shown in this figure, we performed a series of RANS solutions for the full Common Research Model (CRM) aircraft configuration (wing, fuselage, and horizontal tail), which is the same geometry that was used in the Fourth Drag Prediction Workshop (DPW) [64], and it is representative of a long-range transport aircraft. The first row in Fig. 2 shows the friction lines and pressure coefficient, as well as the Lovely–Haines shock sensor (in orange) [65].

To determine if the flow is separated at a given location on the surface, we check if the surface flow velocity has a component in the negative freestream direction (which is approximately the
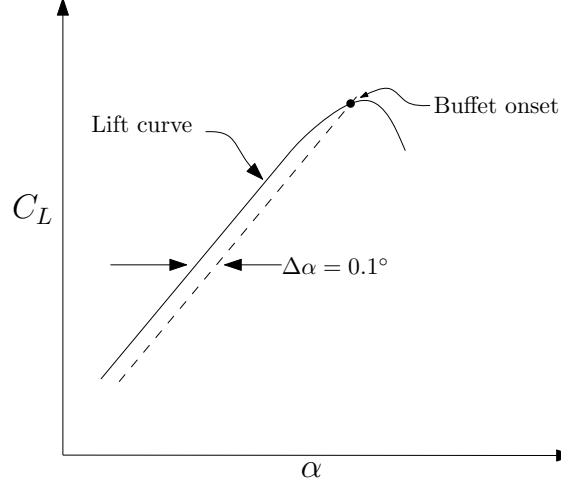
5

Figure 1: Estimating the buffet boundary with the $\Delta\alpha = 0.1$ method.

negative x axis) direction, i.e., if

$$\cos\theta = \frac{\vec{V} \cdot \vec{V}_\infty}{|\vec{V}||\vec{V}_\infty|} < 0, \tag{1}$$

where θ is the angle between the local surface velocity and the freestream. We can then define a *separation sensor* as

$$\chi = \begin{cases} 1 & \text{if} \quad \cos\theta \leq 0 \\ 0 & \text{if} \quad \cos\theta > 0. \end{cases} \tag{2}$$

Thus, χ is specific to each surface location and is a Heaviside function: It is equal to one when the flow is separated, and equal to zero when the flow is attached. The blue areas on the surface for $\alpha = 3.00°$ and $\alpha = 3.29°$ in the bottom row of Fig. 2 show the regions where $\chi = 1$, which approximately coincide with the regions where the flow is separated.
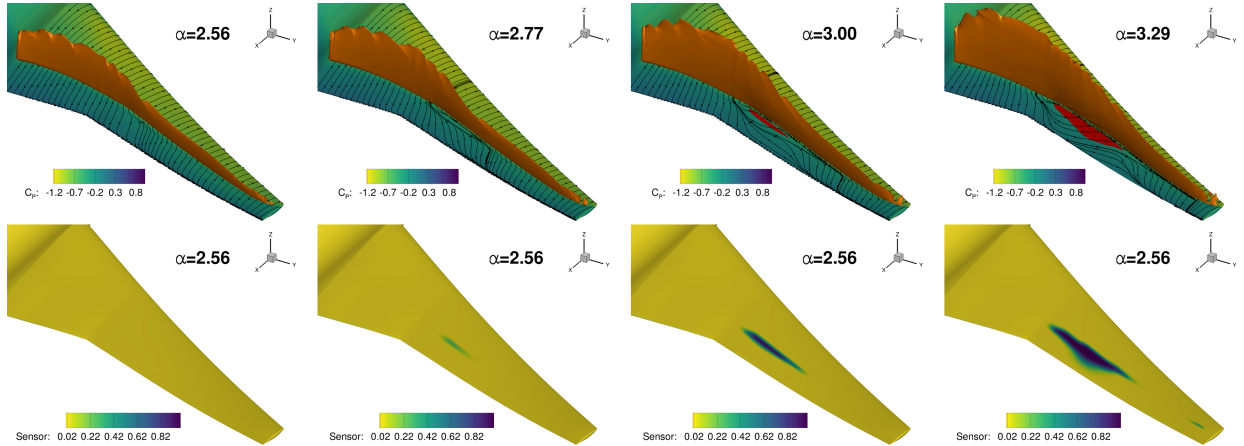

Figure 2: Progression of separated flow for the CRM configuration at $M = 0.85$ with increasing angle of attack. Top row shows the surface streamlines and pressure coefficient, as well as the reversed flow (red) and the shock (orange). Bottom row shows the value of the separation area integrand from Eq. (4).

6

Our hypothesis is that the value of the area where $\chi = 1$ correlates with buffet onset, which is given by the integral of χ over the whole surface area of the wing. Since we need to use this function as a constraint in a gradient-based optimization, we would like this function to be smooth. However, because this integral will be discretized based on a CFD surface mesh, and χ is either zero or one for a given cell, the value of this area does not change continuously with the design variables. To address this issue, we use a smooth Heaviside function to blend the discontinuity as follows:

$$\bar{\chi} = \frac{1}{1 + e^{2k(\cos\theta + \lambda)}}. \tag{3}$$

In this equation, k and λ are free parameters, where k determines the sharpness of the transition, and λ is a parameter that can be used to shift the smoothing function to the left or right as a function of the angle. For our cell-centered solver, the values for $\vec{V}$ are taken from the state variables at the cell center immediately adjacent to the wall, since the velocities at the wall are zero when enforcing the no-slip condition. Figure 3 shows smooth Heaviside functions for $\lambda = -0.1$, 0, 1, and $k = 10$. A value of $k = 10$ is used for all results in this paper. The bottom row of Fig. 2 shows the value of the smoothed separation sensor (4) on the wing surface at $M = 0.85$. The smooth Heaviside function smooths out V_x around the separated flow region. This area formulation can also be applied to constrain other undesirable phenomena, including cavitation [66, 67].
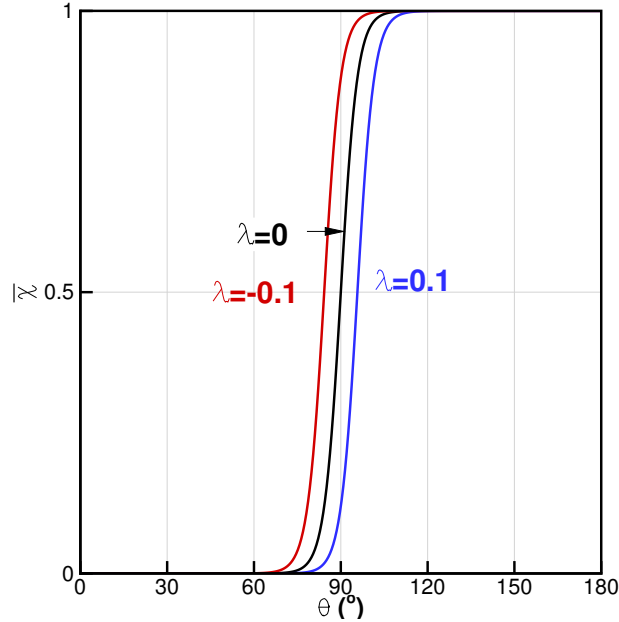

Figure 3: Smoothed separation sensor value (3) versus angle of surface flow. With $k = 10$, the smooth transition occurs over $\pm 15°$.

Next, we can integrate the smooth separation sensor (3) over the surface and normalize it by the aircraft reference area to obtain the proposed *separation metric*:

$$S_{\text{sep}} = \frac{1}{S_{\text{ref}}} \int_S \bar{\chi} \, dS. \tag{4}$$

This is equivalent to performing a weighted area integration of the sensor value shown in the bottom row of Fig. 2.

To find out if the separation metric (4) is correlated to buffet onset, we used the $\Delta\alpha = 0.1$ method as a reference. We start by using the $\Delta\alpha = 0.1$ method to compute the buffet boundary

for the baseline CRM configuration [64] at a flight altitude of 37 000 ft and for Mach numbers ranging from 0.8 to 0.9. The resulting reference buffet boundary is shown as the orange line in Fig. 4. We then plot the lines corresponding to various values of the separation metric, and determine that a cutoff value of $\chi = 4\%$ yields the best agreement when compared with the $\Delta\alpha = 0.1$ method.
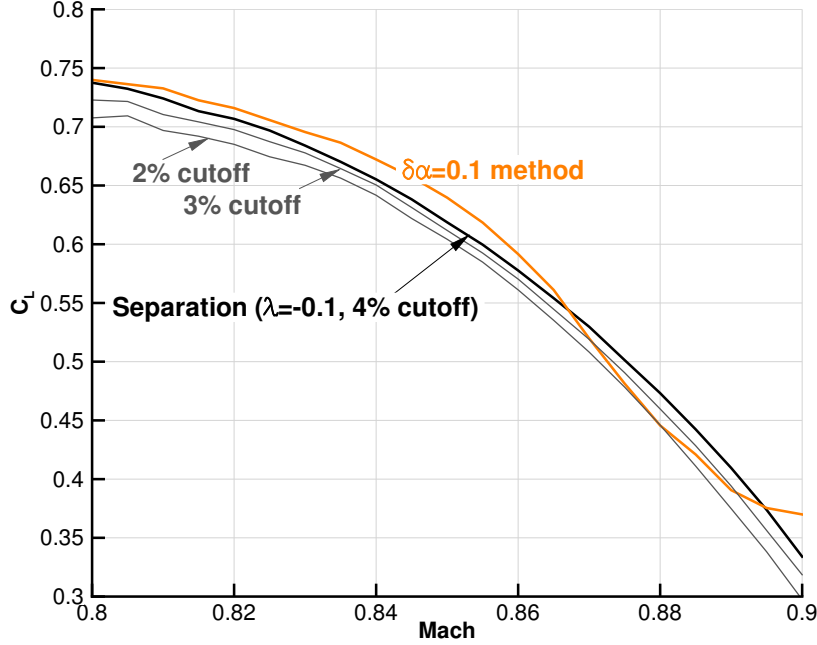

Figure 4: Results of the $\Delta\alpha = 0.1$ method compared with those of the proposed separation metric method.

The overall shape of the buffet onset is consistent with the separation-based criteria, although some discrepancy exists. The discrepancy can be explained by analyzing the slopes of the lift curves for the baseline configuration, which we plot in Fig. 5. The figure also shows the application of the $\Delta\alpha = 0.1$ method, where the line slopes are based on the first two analysis points at the lowest C_L values for each Mach number. The lift curves are close to linear, but the slope of the lift curve increases for Mach numbers in the range 0.81–0.86. Therefore, at these Mach numbers, the $\Delta\alpha = 0.1$ method overpredicts buffet onset compared with the separation-based method, because the slope of the lift curve must decrease more to intersect the offset line. For Mach numbers greater than 0.86, the opposite effect occurs: The lift curves exhibit a reduced slope and intersect the linear offset at lower C_L values, underpredicting the buffet onset. The differences in Fig. 4 are consistent with this effect. As we can see, the separation-based approach varies more smoothly with respect to Mach number, which is beneficial for gradient-based design optimization. Because the separation-based approach is more representative of the actual physics, we believe that it is the more accurate of the two methods.

Figure 6 shows how the separation metric approach is used to construct the buffet-onset boundary. For each α sweep at a fixed Mach number, the intersection of the separation curve with the specified cutoff value determines the boundary. With nearly fully attached flow on the wing upper surface, the sensor metric is close to zero. Next, as α increases, the sensor metric rapidly increases as shock-induced flow separation becomes more severe. The large slope of the separation sensor curve means that the C_L value predicted for buffet is not particularly sensitive to the selected cutoff value.
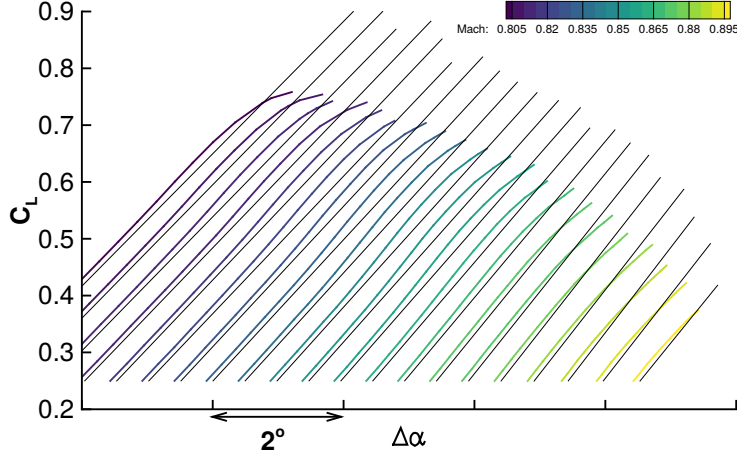
Figure 5: Baseline configuration lift curves and $\Delta\alpha = 0.1$ linear offset for Mach 0.8–0.9. Successive lift curves are offset by $0.5°$.
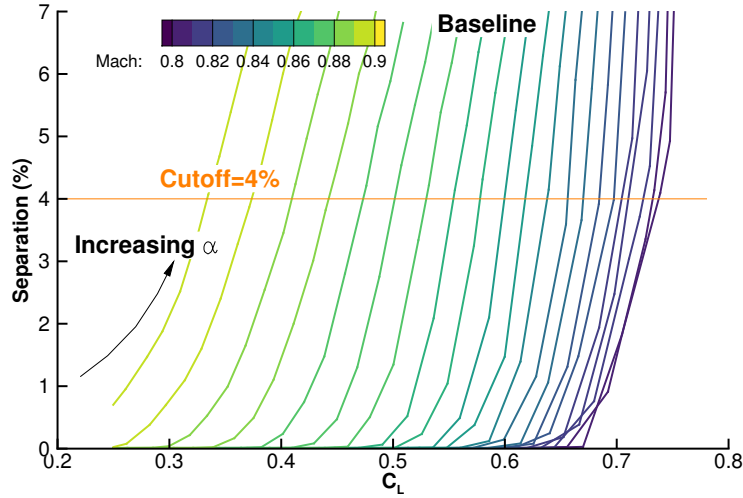

Figure 6: Separation sensor curves for α sweeps and for a range of Mach numbers. The cutoff value indicates the estimated buffet boundary.

To validate the proposed separation metric as a way buffet-onset constraint, we compare the results obtained by using this approach with experimental results by Balakrishna and Acheson [68], who tested the CRM wind tunnel model. They estimated the buffet onset by making high-speed measurements of the strain at the wing root. Buffet onset can be identified by the increase in the strain gauge signal amplitude, which is caused by the shock oscillations interacting with the separated flow. Based on this increase in signal amplitude, they define the *buffet coefficient*, C_B.

Figure 7, which is reproduced from Fig. 4 in Balakrishna and Acheson [68], shows the evolution of the buffet coefficient for two Mach numbers: a high subsonic Mach number ($M = 0.70$) and a transonic Mach number ($M = 0.85$). We overlay lines at α values for which our method yields a separation sensor value of 4%, with $\lambda = -0.1$. We see that the results of the separation sensor method correlate well with the increase in C_B, providing more evidence that the separation-metric approach correctly predicts buffet onset.

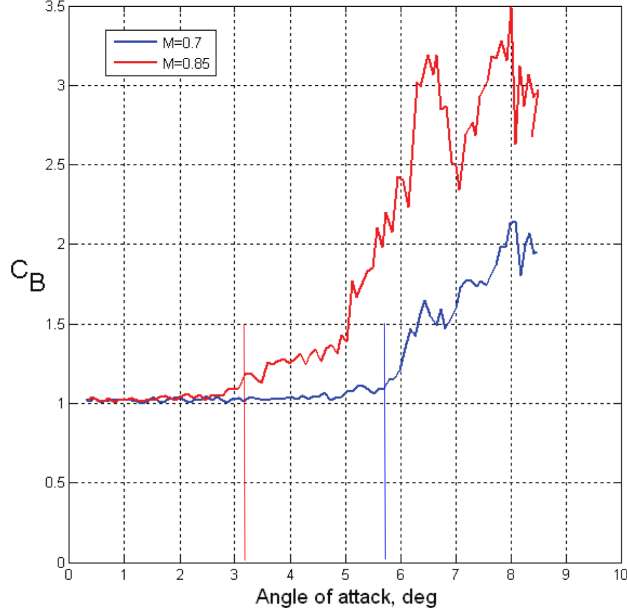Figure 8 shows the surface distribution of C_p and the smoothed separation sensor for the two

Figure 7: Buffet coefficient C_B obtained from wind tunnel data [68]. Vertical lines are the buffet-onset locations predicted by the separation sensor.

flow conditions. Due to the differing freestream Mach numbers, the types of separation and the corresponding separation locations are quite different. For the $M = 0.70$ case, a separation bubble appears just aft of the small leading-edge sonic region, whereas for the $M = 0.85$ case, the separated flow appears near the mid-chord position at the foot of a strong normal shock. Although more comparison with experimental results are required, the separation sensor adequately predicts buffet onset at both high subsonic and transonic flow conditions for this aircraft geometry.

4 Full Configuration Aerodynamic Shape Optimization Benchmark

We now demonstrate the need to consider buffet-onset criteria and the effectiveness of the proposed approach for transonic aerodynamic shape optimization by solving a series of aerodynamic design optimization problems based on the AIAA ADODG Case 5 benchmark [21].

1 Baseline Geometry

The baseline geometry defined in ADODG Case 5 is taken directly from the Fourth Drag Prediction Workshops "Wing-Body-Tail $i_H = 0$" aircraft configuration [69]. This configuration is known as the Common Research Model (CRM) and is representative of a twin-aisle long-range transport. The main reference parameters for the CRM are listed in Table 1.

Quantity	Value
Reference area	$594\,720.0\,\text{in}^2$
Reference chord	$275.8\,\text{in}$
Moment reference	$(1325.90,\ 0,\ 177.95)\,\text{in}$
Reynolds number ($M = 0.85$)	43×10^6

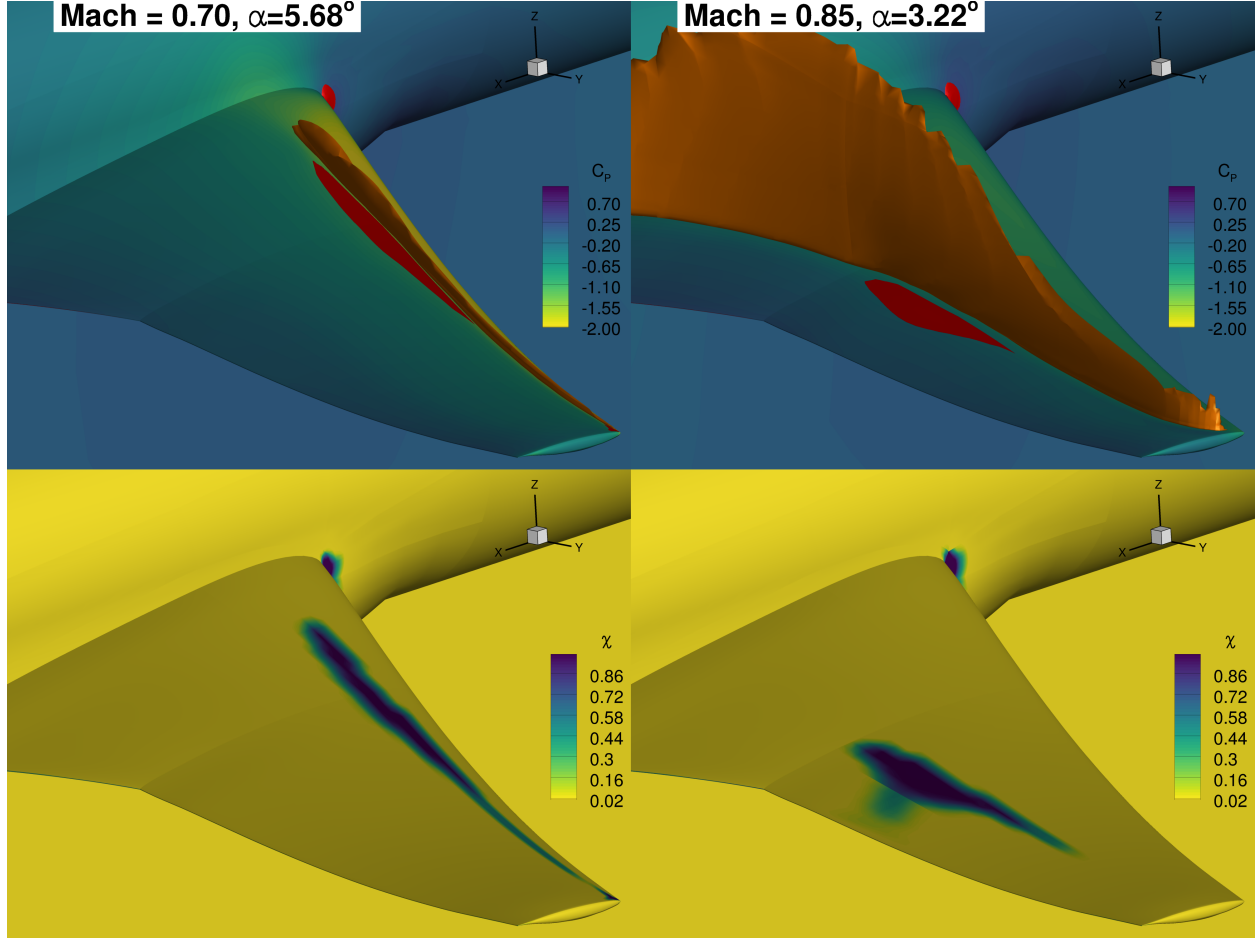Table 1: Reference quantities for CRM full configuration [69].

Figure 8: C_p and smoothed separation sensor surface distribution for the two experimental buffet-onset conditions. In both cases, the separation (red) appears behind the shock (orange).

2 Computational Meshes

We generated a sequence of four CFD meshes for the CRM wing-body-tail configuration using the meshing software ICEM CFD. The meshes are divided into two families: the "1 series" and the "0.5 series." Each "0.5 series" mesh has approximately 2.5 times more cells than the corresponding coarser "1 series" mesh below it and has approximately 3.3 times fewer cells than the next finer "1 series" mesh above it. The two coarsest grids—L2 and L1.5—are used for optimization, whereas the two finest grids—L1 and L0.5—are used only for post-optimization verification purposes. The mesh metrics are summarized in Table 2. Figure 9 compares the surface mesh resolution of the four meshes. Grid convergence studies for the baseline mesh and all optimized configuration meshes are presented in Section 3.5.

3 Optimization Problem Statement

A sequence of seven design optimizations are solved to study the aerodynamic shape optimization of the ADODG full CRM configuration, and to demonstrate the effectiveness of the proposed approach for satisfying buffet requirements. These cases—numbered 5.1 through 5.7—are summarized in Table 3. Only cases 5.1 and 5.2 are currently specified by the ADODG [21]. We added the other cases (cases 5.3 through 5.7) to further study the effects of including buffet-onset conditions

Mesh level	Chordwise cells	Spanwise cells	$y^+_{\max}$	Total cells	Baseline C_D (counts)
L0.5	224	144	~0.5	14 233 600	231.15
L1	168	108	~0.7	5 921 536	234.87
L1.5	112	72	~1.1	1 779 200	249.47
L2	84	54	~1.7	740 192	269.76

Table 2: Mesh characteristics and corresponding trimmed drag coefficient for baseline configuration.
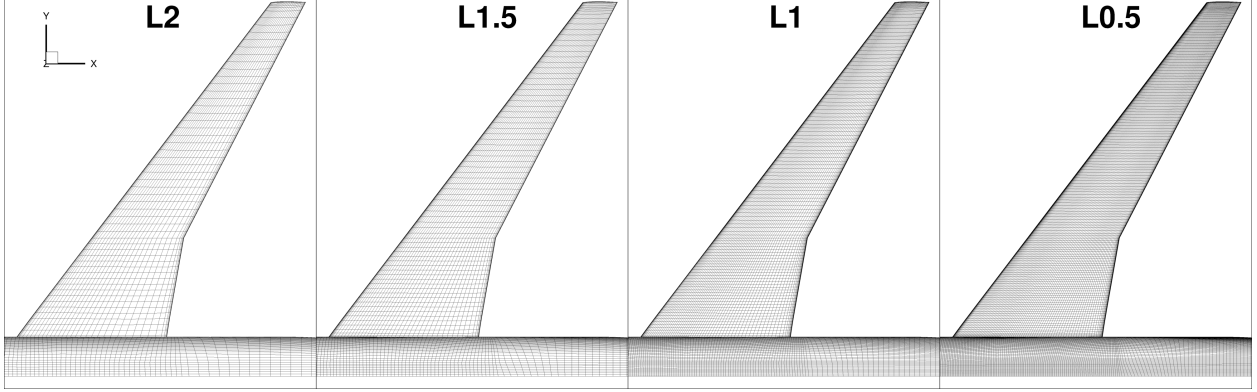

Figure 9: Spatial resolution for each mesh size.

on the optimized geometries. The objective of all optimizations (except for Case 5.7, which is discussed separately) is to reduce the weighted drag coefficient at the N operating conditions. The optimization problem statement can be written as

$$
\begin{array}{lll}
\text{minimize} & \sum_{i=1}^{N} \mathcal{W}_i C_{D_i} & \text{Quantity} \\
\text{with respect to} & \text{Wing cross-sectional shape} & 240 \\
& \text{Wing twist} & 9 \\
& \text{Angle of attack } (\alpha_i) & N \\
& \text{Tail rotation angle } (\eta_i) & N \\
\text{subject to} & C_{L_i} - C_{L_i}^* = 0.0 & N \\
& C_{M_{y_i}} = 0.0 & N \\
& t_j \geq t_{j_{\mathrm{CRM}}} & 750 \\
& S_{\mathrm{sep}_i} \leq 0.04 & N
\end{array}
\tag{5}
$$

Each flight condition i is assigned a weight $\mathcal{W}_i$ that specifies to what extent the drag of the given flight condition influences the objective function. The lift and moment coefficient constraints ensure that the aircraft is trimmed at each flight condition, which can be achieved by the appropriate combination of angle of attack and tail rotation angle. The thickness t_j is computed at 750 points arranged in a 25×30 regular grid in the chordwise and spanwise directions, respectively. These thicknesses are constrained to be greater than or equal to the original thicknesses of the CRM geometry at the corresponding points. Because making the wing as thin as possible is desirable in transonic flow [14], these constraints ensure that the wing does not become too thin, which would result in a significant increase in structural weight. Imposing thickness constraints means that only changes in the wing camber are available to the optimizer.

Only cases 5.4, 5.5, and 5.7 use the separation constraint to satisfy the buffet margin. In these cases, the separation constraint is only enforced for the last two flight conditions, and the

drag coefficient for these conditions does not contribute to the objective function (i.e., $W_j = 0$). Therefore, the adjoint for C_D is not evaluated for the buffet-onset conditions. Conversely, the separation metric adjoint is not evaluated for the conditions where the drag coefficient weights are nonzero. This results in a total of three adjoint solutions being required for for both the cruise and buffet flight conditions, which is desirable from a computational load-balancing perspective.

The ADODG specification for Case 5 disallows the parameterization from modify the planform, and any shape modification may only be made in the vertical direction. Additionally, twist rotation is permitted for the wing, as is a solid-body rotation of the horizontal tail for trimming the aircraft. We use the FFD approach described in Section 2. The FFD volume and the associated geometric design variables are shown in Fig. 10.
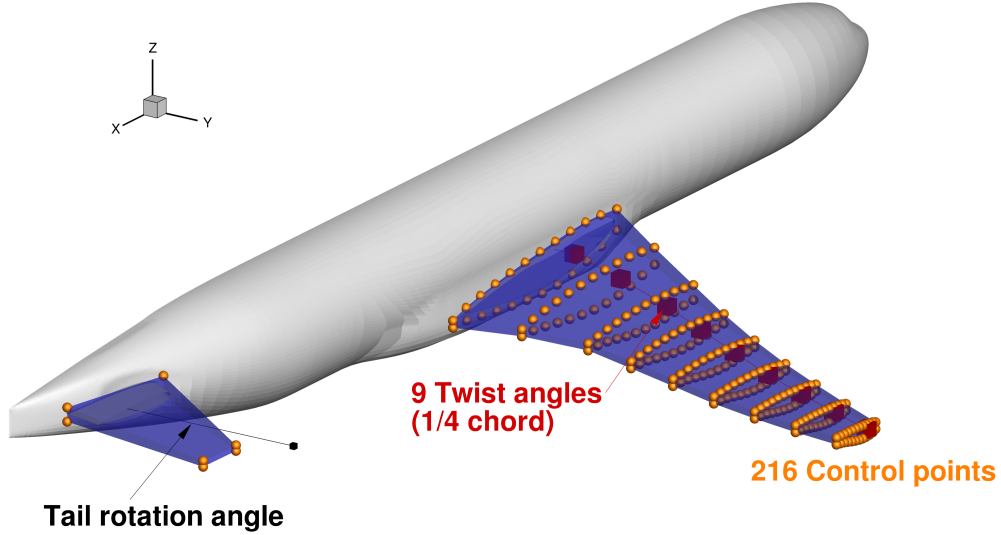

Figure 10: CRM configuration showing the design variables for shape, twist, and tail rotation.

5   Results

1   Multilevel Approach

To reduce the overall computational cost of performing the optimizations, we employ the multilevel optimization approach described previously by the authors [12, 14]. Optimizations are first carried out on the coarsest mesh (L2), and the resulting optimum design becomes the starting design for the next finer mesh (L1.5), and so on. In this work, only the first two grid levels are used for the optimization. Because optimizing on the coarse grid costs less, we can afford to do more iterations on this grid. For this approach to be effective, the coarse grid must capture the main characteristics of the flow.

Figure 11 compares the baseline and optimized designs for Case 5.1. The aircraft planform views show the baseline and optimized designs obtained by using the L2 grid, and the L1.5 optimization obtained by using the L2 optimized shape as the starting point. Color-coded slices of the airfoil shapes and the corresponding C_p distributions are shown for four spanwise locations at the bottom of the figure. We see that the coarse optimization (using the L2 grid) successfully eliminates the shock on the upper wing surface, resulting in parallel isobars. Even without further optimization, almost all of the drag improvement predicted by the coarse grid is realized on the fine grid. A comparison of the orange and black lines on the outer C_p distributions shows that the only significant difference is the appearance of a weak shock on the refined grid. The fine optimization further
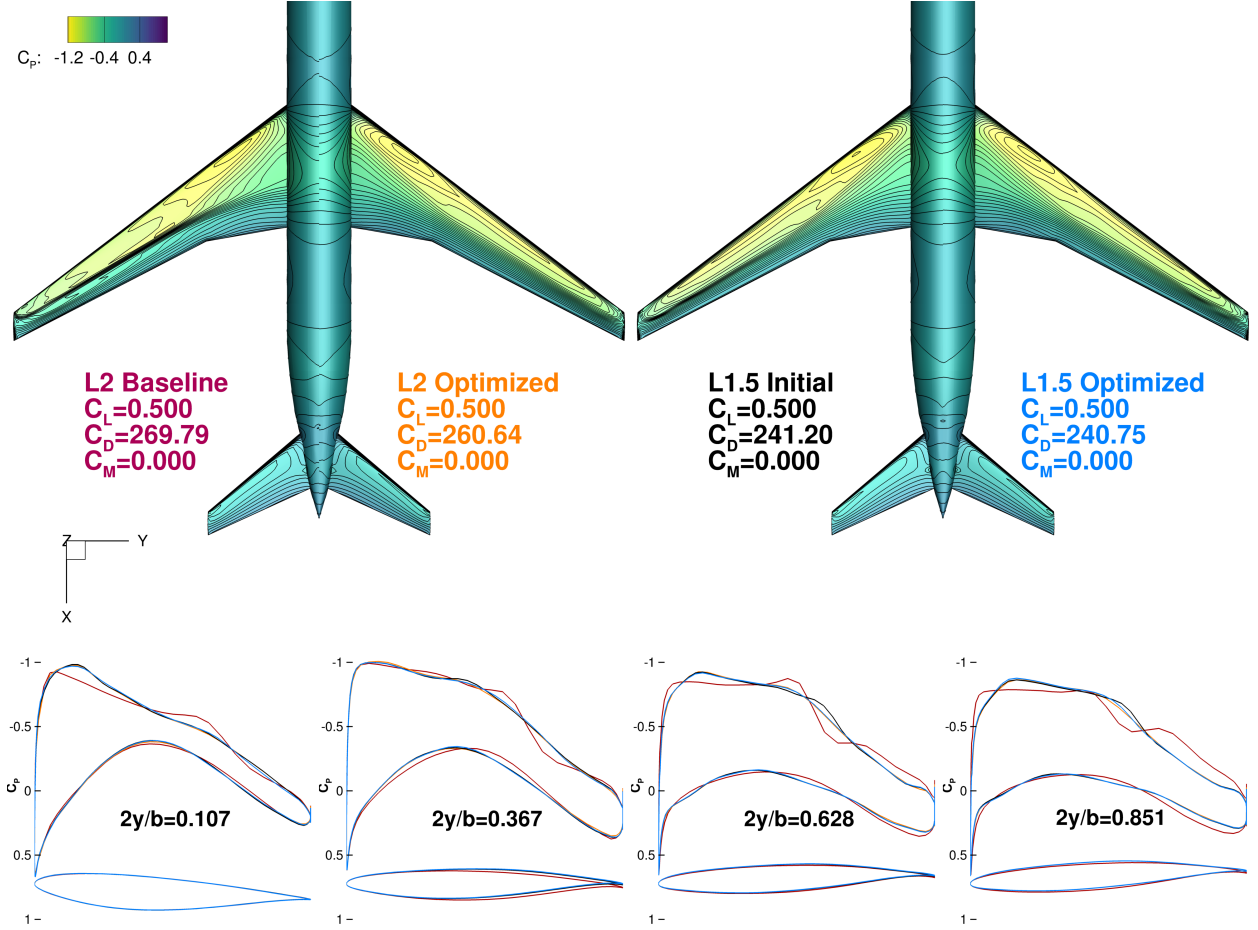
Figure 11: Baseline design compared with optimized designs for Case 5.1. The coarse-grid optimum is a good starting point for the fine-grid optimization.

improves the design, eliminating this shock and lowering the drag even further. This behavior is consistent with our previous results, where we used three uniformly refined grid levels [12]. We employ this multilevel approach for all optimizations in the present work.

2   Optimization Results

In this section we present the main results for each CRM aerodynamic shape optimization (cases 5.1 through 5.7). Figure 12 shows the evolution of the SNOPT merit function and optimality. The merit function is the value of the augmented Lagrangian given by SNOPT, which becomes the same as the objective function value once all the constraints are satisfied toward the end of the optimizations. The optimality is the residual of the Karush–Kuhn–Tucker (KKT) optimality conditions, which measures how well the optimization has converged [57].

The optimality tolerance was set to 10^{-5}, which is achieved for most optimizations. The L2 optimizations are limited to 150 iterations, whereas the L1.5 optimizations are limited to a further 50 iterations. Case 5.7 uses a different objective function and is a maximization instead of a minimization. Generally, the finer optimizations with the L1.5 mesh achieved the same convergence tolerance as the coarse L2 mesh.

Figures 13 and 14 summarize the key features of the two ADODG optimizations (cases 5.1
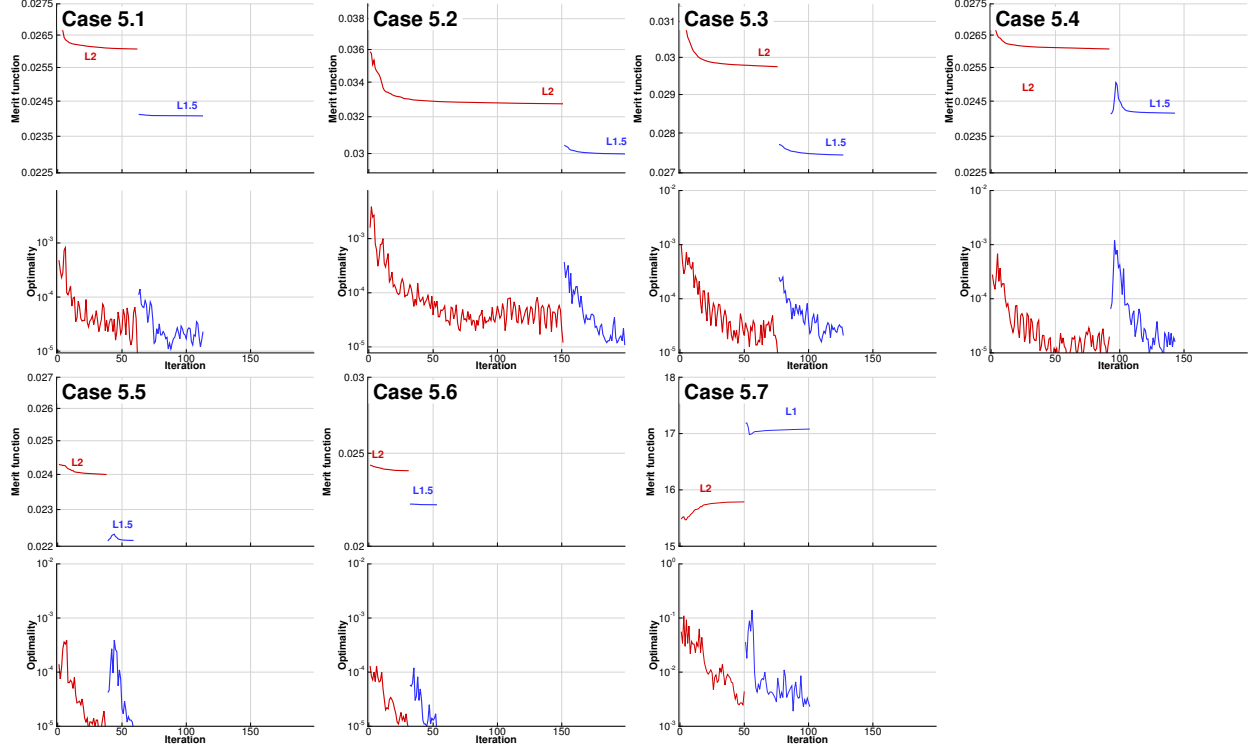
Figure 12: Merit function and optimality evolution for each optimization case.

and 5.2, respectively). The results of the baseline configuration are shown in red, whereas the optimized results are shown in blue. The planform view of the wing and fuselage shows the C_p contours of the baseline geometry (left) and the optimized geometry (right) under the nominal operating conditions ($M = 0.85$, $C_L = 0.5$). Just below the planform view, the front view also shows the C_p contours and adds a visualization of the shock surface [70]. Below the front view, we plot the spanwise distributions of lift, twist, and thickness-to-chord ratio (t/c). A reference elliptic lift distribution is shown in gray. The right side of the figure displays the cross-sectional shapes and C_p distributions at the five spanwise locations indicated by the labels A–E in the planform view. Finally, the bottom-right plot shows the drag divergence behavior for three lift coefficients: $C_L = (0.45, 0.50, 0.55)$.

The single-point optimization (Case 5.1) is similar to the wing-alone optimization done in previous work [12], where it is referred to as "aerodynamic shape optimization without thickness reduction." In that case, the wing-alone configuration was optimized at the same Mach number and lift coefficient, but at a much lower wind-tunnel Reynolds number of 5×10^6. The previous work indicates that a 10.5 drag count reduction is possible for the CRM wing-alone configuration. This compares well with the 8.6 count reduction that we obtain in the full wing-body-tail configuration studied herein. The cross-sectional plots of the airfoils at various spanwise sections show how little the shape needs to be modified to obtain a substantial change in performance. The drag divergence curves highlight the single-point design nature of the optimized configuration. A drag dip is present at the on-design condition, but the performance is worse at most other Mach numbers and lift coefficients.

When we introduce the $\alpha = 0.1$ method above for predicting buffet onset, we show the lift curves for the baseline configuration in Fig. 5. We now compute the same curves for the optimized configuration of Case 5.1, as shown in Fig. 15. The deviation from the linear slope observed in the
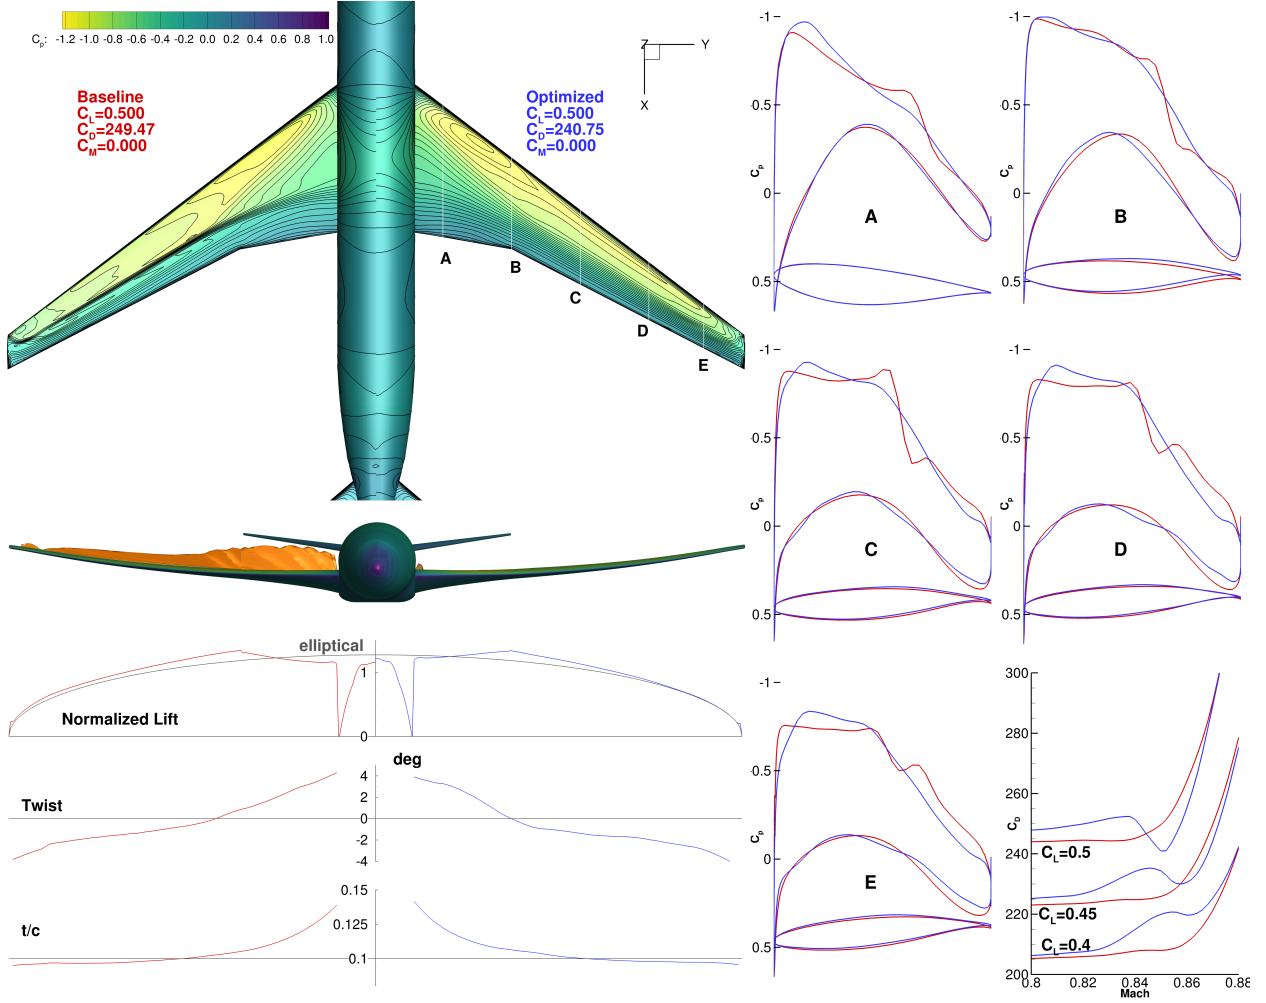
15

Figure 13: Baseline analysis (red curves) compared to the Case 5.1 result (blue curves). This single-point optimization leads to high performance at the nominal operating condition.

baseline is more pronounced for this optimized aircraft, which means that including a physics-based buffet-onset constraint as we propose is all-the-more crucial; otherwise the optimizer will exploit the lack of such a constraint and produce designs that are not realizable.

Case 5.2 adds two additional equally weighted operating conditions near the buffet-onset boundary. Unlike Case 5.1, for which we obtain a shock-free wing, Case 5.2 results in double shocks at the nominal operating condition. In this case, the drag at the nominal operating condition actually increases by 2.8 counts, as shown in Fig. 14. The drag divergence curves indicate a significant drag penalty across the lower Mach numbers, but this design does have a much higher drag divergence Mach number than the baseline design.

Although drag-coefficient divergence curves yield useful insights into optimized designs, examining the performance in the full M-C_L space is particularly instructive. In the context of transonic transport wing design, ML/D is a better measure of performance because it includes the benefit imparted on overall aircraft efficiency by a higher cruise speed. This overall performance can be approximated by the Breguet range equation

$$R = \frac{Ma}{c}\frac{L}{D}\ln\left(\frac{W_1}{W_2}\right), \tag{6}$$
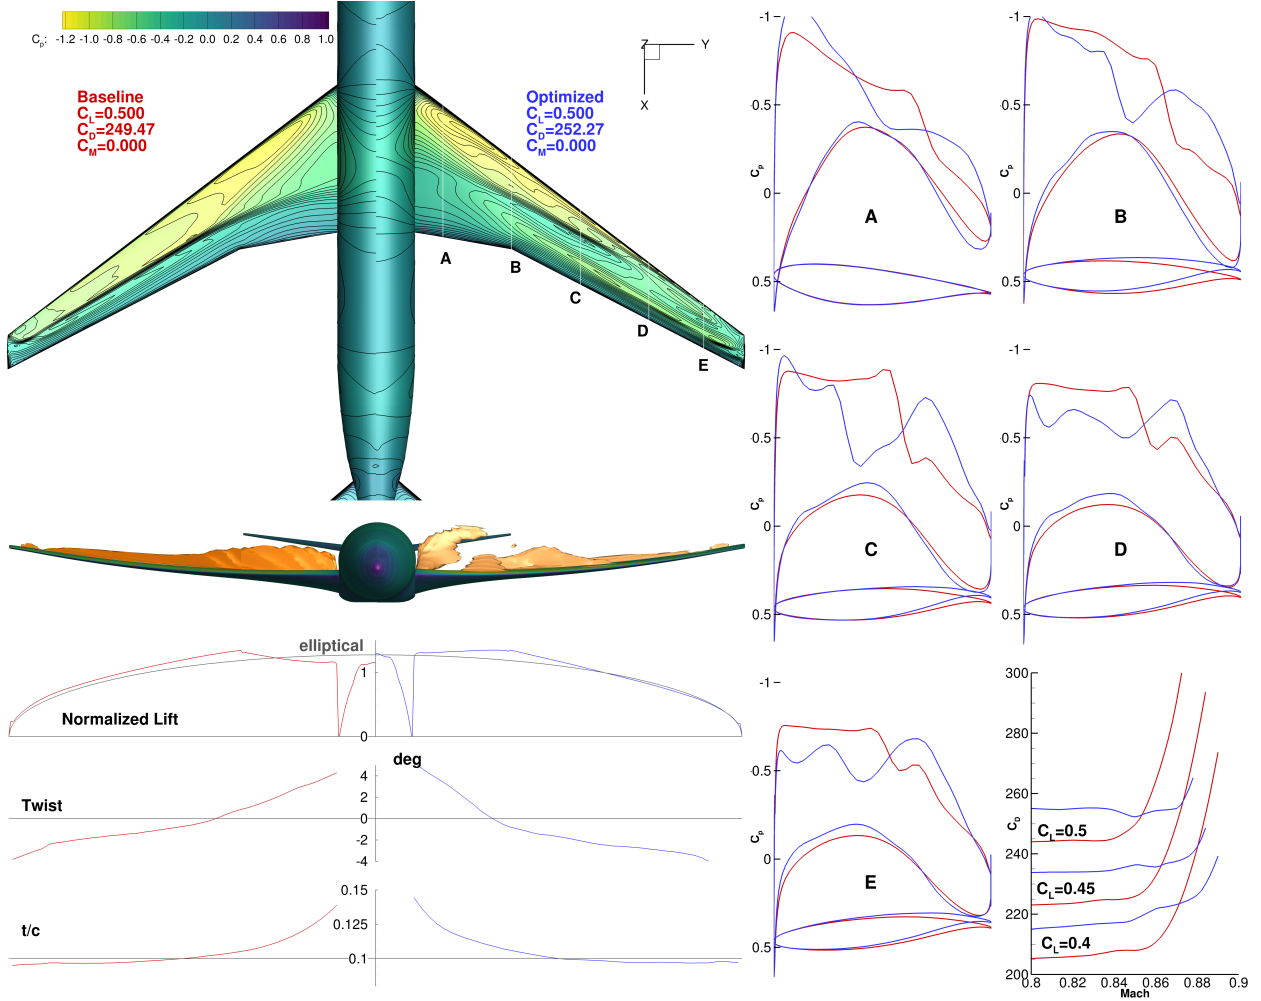
16

Figure 14: Baseline analysis (red curves) compared to the Case 5.2 result (blue curves). To obtain a small improvement at the highest Mach numbers, performance is sacrificed across a large range of Mach numbers.

where L/D is the lift-to-drag ratio, a is the speed of sound, c is the thrust-specific fuel consumption, and W_1 and W_2 are the initial and final cruise weights, respectively. For a purely aerodynamic optimization at fixed Mach number, only L/D varies if we assume a constant c and weight ratio W_1/W_2, so we are left with ML/D.

The procedure for generating contour plots is detailed in Appendix A. The contour plots are generated by using the L1.5 grid, and we ignore the additional drag associated with the nacelle, pylon, and vertical stabilizer. Figure 16 shows contour plots for all seven optimizations and the baseline design.

The contours in each figure extend up to the predicted buffet-onset curve shown in red. The orange curve shows the buffet onset predicted by using the $\Delta\alpha = 0.1$ method described in Section 3. Several regions appear where the orange curves are missing data, which we attribute to the failure to find an intersection between the two lift curves. Overall, the separation-metric method continues to produce results that are close to those of the $\Delta\alpha = 0.1$ method, despite the large changes in the buffet-onset boundary.

The blue curve represents the 30% margin to buffet-onset boundary and is computed directly
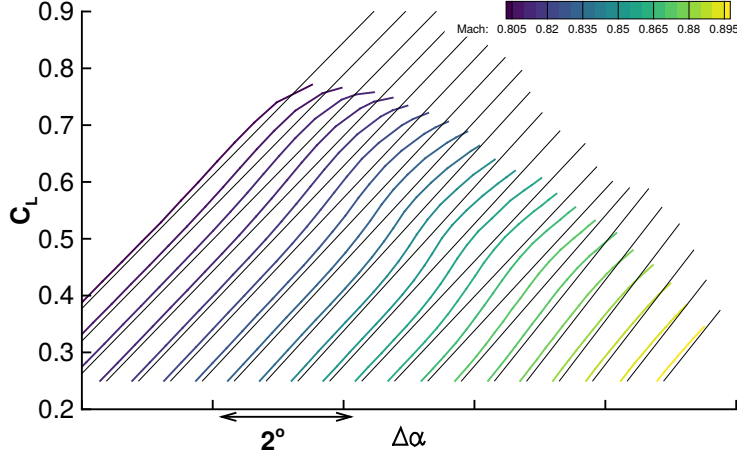
17

Figure 15: Lift curves for the Case 5.1 configuration with successive lift curves offset by 0.5°. The lift curves for the optimized configuration are more nonlinear.

from the red buffet-onset curve. For normal operation, only operating conditions below the buffet-margin curve can be considered. The absolute maximum ML/D value for each configuration is shown in pink. Two specific contours for the optimization configuration (one for the baseline configuration) are highlighted: the contour of 99% $(ML/D)_{\max}$ for the particular design is shown in blue, and the contour of 99% $(ML/D)_{\max}$ for the baseline configuration is shown in red. The motivation for plotting these 99% contours is that airliners typically fly between the Mach number yielding maximum range (approximated by the maximum ML/D value in the figures) and a higher Mach number that yields a 1% fuel-burn penalty but decreases in the flight time. The area enclosed by both of these contours is used to quantify the robustness of the design in these figures. The areas are scaled by a factor of 100^2 so that the area of the rectangle measuring 0.01 in M and 0.01 in C_L has unit area.

The design operating conditions listed in Table 3 are shown as diamonds. The operating conditions considered for the objective function are shown in black, whereas the buffet-onset constraint conditions are shown in red. The first buffet point ($M = 0.85$, $C_L = 0.65$) is at the nominal cruise Mach number and the C_L value corresponding to a 1.3g maneuver. The second buffet point ($M = 0.89, C_L = 0.456$) is 0.04 higher in Mach number, which is a typical margin between a nominal cruise Mach number and the maximum Mach number (M_{MO}) condition. The lift coefficient for this condition is adjusted to give the same dimensional lift as the nominal cruise condition at the same altitude.

Two additional regions are highlighted in black and orange, which we refer to as *integration regions*. They are constructed as follows: The Mach range is from 0.83 to 0.86, which corresponds to the typical range of operating Mach numbers for an aircraft such as the CRM. The upper line corresponds to the buffet-margin boundary, which is equivalent to specifying the maximum altitude the aircraft can fly for a particular weight. The bottom line corresponds to the reduced C_L for a 4000 ft decrease in altitude. To put it in another way, the integration region contains all operating conditions within 4000 ft of the buffet-constrained ceiling and at all normal operating Mach numbers. It is likely that the aircraft spends the vast majority of cruising flight in this region. The black integration region corresponds to the baseline design, whereas the orange regions are adjusted to reflect the actual buffet-margin boundary for each design. In addition, the upper edge of the black region indicates how the buffet-onset boundary changes for each design relative

18

to the baseline configuration for the specific Mach range of integration.

Figure 17 displays a different visualization of the data already shown in Fig. 16. Here, we plot the percent change of each design relative to the baseline configuration. Note that the plot region is limited to flight conditions below the buffet boundary corresponding to each design. The color of the boundary indicates which is one active: The black boundary indicates the baseline buffet boundary is active, meaning that the optimized design boundary is higher than the baseline. The orange boundary means that the optimized boundary is active and thus lower than the baseline buffet boundary. The integration region for each configuration is also shown in orange.

The contour plots give a much more complete understanding of the optimized designs. Unsurprisingly, the single-point optimization—Case 5.1—produces the highest ML/D value, which is almost exactly matched to the design operating condition. However, with no way to constrain the buffet-onset boundary, the value of $(ML/D)_{\max}$ is now above the buffet-margin boundary, which means that this high-performance point cannot be achieved in practice because it falls outside the normal flight envelope. The 99% $(ML/D)_{\max}$ contour (blue) is small, indicating a highly localized point design. Despite the high ML/D value, the average ML/D in its own integration region (orange) is 4.6% worse than that for the baseline design.

For Case 5.2, the addition of operating conditions at the edge of the buffet-onset envelope substantially improves the buffet boundary over the entire range of Mach numbers. This case results in the most robust buffet-onset behavior of all cases. However, the value of $(ML/D)_{\max}$ barely improves over that of the baseline design (17.18 vs 17.13). Worse still, as in Case 5.1, the high-performance region lies almost entirely outside the buffet-margin boundary, rendering the high-performance region unattainable. Even for this case, the average ML/D in the integration region is slightly worse (-0.5%) than for the baseline design. Note that the increased performance afforded by the higher buffet boundary is only possible if the baseline aircraft is buffet limited in altitude over the specific range of Mach numbers, as opposed to thrust limited. If the aircraft were thrust limited over the integration range, the obtainable performance would be the integral over the black integration region..

In Case 5.3, we attempt to improve upon Case 5.2 by reducing the weighting factor for the near-buffet conditions. For this case, the nominal operating condition has a weight of 2/3, whereas the remaining two points each have weights of 1/6. The adjusted weightings yield a much more useful design. This is the first case where a significant portion of the 99% $(ML/D)_{\max}$ contour falls within the integration region. In addition, the design is robust, as evidenced by the larger area enclosed by the blue contour when compared to the baseline design. As with the two previous cases, the increased performance is only possible if the aircraft can operate at higher altitudes. The other problem with this case is that the specific weightings are picked arbitrarily. Although these particular weights yield acceptable results, these weight values are not guaranteed to work well for another configuration or optimization problem.

Case 5.4 is the first optimization to use the separation sensor directly as an optimization constraint. Case 5.4 retains the same operating conditions as cases 5.2 and 5.3, but instead of having the drag from the flight conditions near the buffet boundary contribute to the drag objective function, it uses the buffet-onset flight conditions to compute the separation sensor and constrain its value. Note that a slight discrepancy exists between the operating conditions (red diamonds) and the buffet-onset boundary itself. The reason for this result is that the buffet-onset conditions are analyzed by using the scalar Jameson–Schmidt–Turkel dissipation scheme [71], which results in a solution with more dissipation than the matrix dissipation scheme. The scalar scheme provides the increased robustness necessary for the optimization, which is unnecessary for the contour plot evaluations. The more dissipative scalar scheme slightly underpredicts the area of separated flow, so the buffet boundary in the contour plot is lower when analyzed with the matrix scheme for the

contour plot. Overall, the performance of this design is similar to that obtained with single-point optimization (Case 5.1). Most of the high-performance region lies outside the integration region. However, the reduction in performance is not as pronounced as with Case 5.1, for which performance is reduced by just 1.6% in the original integration region and by almost zero in the on-design integration region. Nevertheless, there is a small improvement in the buffet-onset boundary.

Upon analyzing the results from cases 5.1–5.4, we noticed that the nominal design point always appears toward the upper side or even completely outside the integration region, and that all optimization discussed thus far fails to improve the performance in the baseline integration region. To address this issue, we formulate a multipoint optimization (Case 5.5). Previous optimizations performed by the authors on the CRM wing-alone configuration show that selecting five operating conditions arranged as a cross in the M-C_L space results in highly robust designs [14]. Because our goal is to improve the performance in the original integration region, we distribute the five conditions as follows. The nominal Mach number is reduced to 0.845 for the first three operating conditions. The first point lies on the $1.3g$ buffet-onset boundary, whereas the next two points are at C_L values corresponding to 2000 and 4000 ft below. The two remaining points are 2000 ft lower than the buffet-onset boundary with a variation in Mach number of ± 0.01 . The buffet-onset conditions are taken at $M = 0.85$, $C_L = 0.65$, and $M = 0.89$, $C_L = 0.41$. The latter point is taken from the baseline design buffet-onset boundary. The overall performance of this case is superior to all previously discussed cases. The performance in the baseline integration region increases by 1.2%, and the performance of the updated integration region increases by 2.0%. The design is very robust, as shown by the area inside the 99% $(ML/D)_{\mathrm{max}}$ contour. Furthermore, the point of maximum performance appears inside the operating envelope. Given these results, trying a lower nominal C_L for cases 5.2–5.4 may be worth considering in the future.

Next we developed Case 5.6, which is designed to investigate the effect of removing the buffet-onset conditions present in Case 5.5. We wish to answer the following question: Is a multipoint optimization near the design operating condition sufficient to ensure a robust buffet-onset envelope? Unsurprisingly, without the buffet-onset constraints, the buffet margin boundary drops slightly over the integration envelope, pushing the integration region into a lower-performance region. The average value of ML/D for the integration region is 16.88, only 0.6% higher than for the baseline configuration and much smaller than the 2.0% improvement obtained in Case 5.5.

Finally, for Case 5.7, we formulate a different design optimization problem. We wish to remove the requirement of specifying fixed design lift coefficients and let the optimization itself determine the ideal on-design condition. All cases presented thus far are lift-constrained drag minimizations with fixed operating conditions. The fixed operating conditions also include fixing the value of C_L for the buffet-onset locations. In the formulation of Case 5.7, we want the optimization to directly adjust the single nominal operating condition. The remaining operating conditions are then explicitly linked to this design C_L. More specifically, the high-C_L buffet-onset conditions must have 1.35 times the lift of the nominal cruise Mach. However, to achieve a higher buffet-onset boundary for which more of the integration region lies inside the 99% ML/D_{max} contour, we use a factor slightly greater than the minimum factor of 1.3. The high-Mach buffet case must have the same physical lift as the nominal operating condition at $M = 0.89$ and at the same altitude. Finally, the remaining operating conditions move vertically in sync with the changing design C_L.

The modified optimization formulation is as follows:

$$
\begin{array}{llr}
\text{maximize} & \sum_{i=1}^{N} \mathcal{W}_i M_i L_i / D_i & \text{Quantity} \\
\text{with respect to} & \text{Wing cross-sectional shape} & 240 \\
& \text{Wing twist} & 9 \\
& \text{Angle of attack } (\alpha_i) & N \\
& \text{Tail rotation angle } (\eta_i) & N \\
& \text{Design } C_L & 1 \\
\text{subject to} & C_{L_i} - C_{L_i}^{*} = 0.0 & N \\
& C_{M_{y_i}} = 0.0 & N \\
& t_j \geq t_{j_{\mathrm{CRM}}} & 750 \\
& \mathrm{Sep}_i \leq 0.04 & N
\end{array}
\tag{7}
$$

Note that the operating conditions (diamonds) shown in Fig. 16 are the optimized values. The optimization increased the nominal design C_L from the initial value of 0.490 (the value used in Cases 5.5 and 5.6) to 0.520. This increase is made possible by a corresponding increase in the buffet-onset boundary. The previous optimizations, especially Case 5.2, showed that there can be a significant penalty in cruise drag for a higher buffet boundary. For Case 5.7, we have given the optimizer sufficient information to make this trade-off optimal. This results in a slightly higher average performance than Case 5.5 (17.13 vs. 17.11), as well as a higher buffet-onset boundary. The design is also highly robust, exhibiting the largest 99% $(ML/D)_{\mathrm{max}}$ contour of all the cases.

Further insight into the differences between the optimized designs is provided by Fig. 17. It is particularly interesting to see that there is a region between $M = 0.86$ and $M = 0.88$ at low C_L that is universally worse on all optimized designs. This is particularly noticeable on the single-point designs (Cases 5.1 and 5.4). It is least evident in Case 5.5, where there is an improvement over almost the entire contour region. Compared to Case 5.7, the higher buffet-onset performance appears to be correlated with the reduced low-C_L performance. The performance reduction at lower lift coefficients in Case 5.7 is limited to less than 2%, which is acceptable given the performance increase at the higher lift coefficients.

3 Grid Convergence

We studied grid convergence for the baseline geometry and for all optimized configurations. For the grid-convergence studies, we apply the optimized geometry from the L1.5 mesh to the each of the four meshes in sequence. The drag convergence for each mesh configuration is shown in Fig. 18. The drag coefficient, when plotted against the grid factor $N_{\mathrm{cells}}^{-2/3}$, is approximately linear, indicating second-order convergence. However, the finest mesh analyzed (L0.5) does not fall directly on the line, which indicates that a more highly resolved mesh is necessary to determine a grid-converged value. However, for aerodynamic shape optimization, we are generally more concerned with the *change* in drag coefficient resulting from a design change as opposed to the grid-converged drag coefficient. Figure 19 shows the change in drag coefficient for each configuration on each mesh level. Remarkably little variation occurs across each mesh level, which we attribute to the fact that the spurious drag remains roughly constant for a given mesh, independent of the design modifications. The maximum variation between the improvement on the L1.5 mesh—the finest mesh used for optimization—and the L0.5 mesh is 0.63 counts (in Case 5.6). Given the much larger computational cost of optimizing with the L1 or L0.5 meshes and the small difference in predicted drag improvement when using these finer meshes, the use the L1.5 mesh for our optimizations is a good choice.

4 Computational Cost

Multipoint three-dimensional RANS-based aerodynamic shape optimizations are costly from a computational perspective, so we make every effort to reduce the total cost of the optimizations. Table 4 lists the total CPU cost, in processor hours, required to generate the results presented in this paper. All the computations were performed on nodes with two 4-core E5540 CPUs running at $2.53\,\text{GHz}$ with $16\,\text{GB}$ of RAM per node. The nodes are connected with QDR InfiniBand. All L2 and L1.5 meshes were run on 64 cores, whereas the L1 and L0.5 meshes for the grid convergence study used 128 cores.

The optimization consumed 63% of the total computational time, and the remainder of the time was used for postprocessing. The ML/D contour plots are particularly costly, because each plot requires approximately 400–500 individual CFD evaluations.

6 Conclusions

We present herein a new formulation for predicting buffet onset and for effectively implementing it as a design optimization constraint. The proposed method is based on the integration of a separation sensor along with a cutoff value to estimate when buffet first occurs and can be directly evaluated with only one steady RANS CFD solution. The results of this method compare well with those of the $\Delta\alpha = 0.1$ method for the CRM configuration, and for various optimized designs. A comparison with experimental data obtained from wind tunnel test also shows that the proposed model has good predictive capabilities. The separation sensor method is particularly well suited for formulating a constraint in gradient-based optimization, because it is easy to implement in a discrete adjoint optimization framework, and the resulting function (although highly nonlinear) is smooth.

To demonstrate the effectiveness of the proposed approach, we optimize seven CRM wing-body-tail configurations. All optimizations are done with respect to 216 shape variables, 9 twist variables, and tail rotation angle, subject to lift, pitching moment, volume, thickness, and separation constraints. To reduce the overall computational cost, we apply a two-level sequential optimization approach. At the nominal operating condition of $M = 0.85$, $C_L = 0.5$, the single-point optimization (Case 5.1) reduces the drag coefficient from 249.5 counts to 240.9 counts, a reduction of 3.4%. For a more complete comparison of the optimized designs, we plot contours of ML/D in M-C_L space, which provides a visual and intuitive way of comparing the performance and robustness of the optimized configurations.

For Case 5.2, we add two operating conditions near buffet onset, which increases substantially the performance at these points and produces a high buffet-onset boundary. However, the overall performance as measured by a typical operating envelope is lower than the baseline. Although weighting the nominal cruise point (Case 5.3) more than the off-design condition improves performance, it requires knowledge of how to choose the appropriate weights, which may be case dependent.

In Case 5.4, we introduce the use of the separation metric to directly control the buffet-onset and buffet-margin boundaries. Although this approach is effective, the overall performance of the optimized design is unsatisfactory; it is lower than that of the weighted-points approach in Case 5.3.

The remaining cases use five main operating conditions to produce more robust designs. Performance improves for Case 5.5 over nearly the entire transonic range, with a simultaneous improvement of part of the buffet-onset boundary. Case 5.6 removes the buffet-onset conditions, demonstrating the insufficiency of a multipoint optimization with all conditions near the on-design condition. In this case, the buffet margin boundary encroaches onto the cruise performance region, reducing the average usable improvement from 2.0% for Case 5.5 to only 0.6% for case 5.6. Finally,
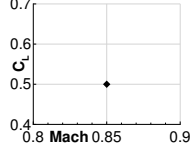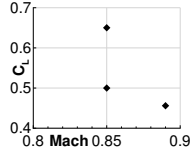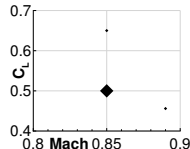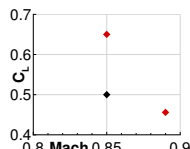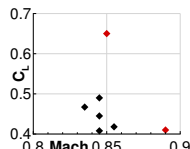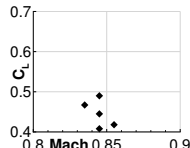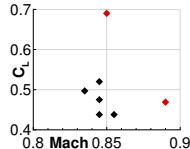
in Case 5.7, we do an ML/D maximization with automatic determination of the operating C_L, which improves the average performance above that of Case 5.5 while pushing the buffet-onset boundary beyond the operational envelope.

Given these results, we recommend that physics-based buffet-onset constraints be enforced for aerodynamic and aerostructural shape optimization of transonic transports, as we do in this work. The separation metric we develop herein is easily implemented and yields robust results, so it provides a much needed constraint formulation for the aerodynamic shape optimization community.

7   Acknowledgments

Table 3: Operating conditions for each optimization. Red diamonds denote separation-constrained points. Operating conditions for Case 5.7 are determined by the optimization process itself. Zero weight means that only the flight condition is considered for the constraints.

| Case	Point	Weights ($\mathcal{W}_i$)	Mach	C_L	Re	M–C_L plot
5.1	1	1	0.85	0.500	43.00×10^6	
5.2	1	1/3	0.85	0.500	43.00×10^6	
	2	1/3	0.85	0.650	43.00×10^6	
	3	1/3	0.89	0.456	45.00×10^6	
5.3	1	2/3	0.85	0.500	43.00×10^6	
	2	1/6	0.85	0.650	43.00×10^6	
	3	1/6	0.89	0.456	45.00×10^6	
5.4	1	1	0.85	0.500	43.00×10^6	
	2	0	0.85	0.650	43.00×10^6	
	3	0	0.89	0.456	45.00×10^6	
5.5	1	1/5	0.845	0.490	42.75×10^6	
	2	1/5	0.845	0.445	42.75×10^6	
	3	1/5	0.845	0.408	42.75×10^6	
	4	1/5	0.835	0.467	42.24×10^6	
	5	1/5	0.855	0.418	43.25×10^6	
	6	0	0.85	0.650	43.00×10^6	
	7	0	0.89	0.456	45.00×10^6	
5.6	1	1/5	0.845	0.490	42.75×10^6	
	2	1/5	0.845	0.445	42.75×10^6	
	3	1/5	0.845	0.408	42.75×10^6	
	4	1/5	0.835	0.467	42.24×10^6	
	5	1/5	0.855	0.418	43.25×10^6	
5.7	1	1/5	0.845	0.520	42.75×10^6	
	2	1/5	0.845	0.475	42.75×10^6	
	3	1/5	0.845	0.438	42.75×10^6	
	4	1/5	0.835	0.497	42.24×10^6	
	5	1/5	0.855	0.448	43.25×10^6	
	6	0	0.85	0.690	43.00×10^6	
	7	0	0.89	0.469	45.00×10^6	
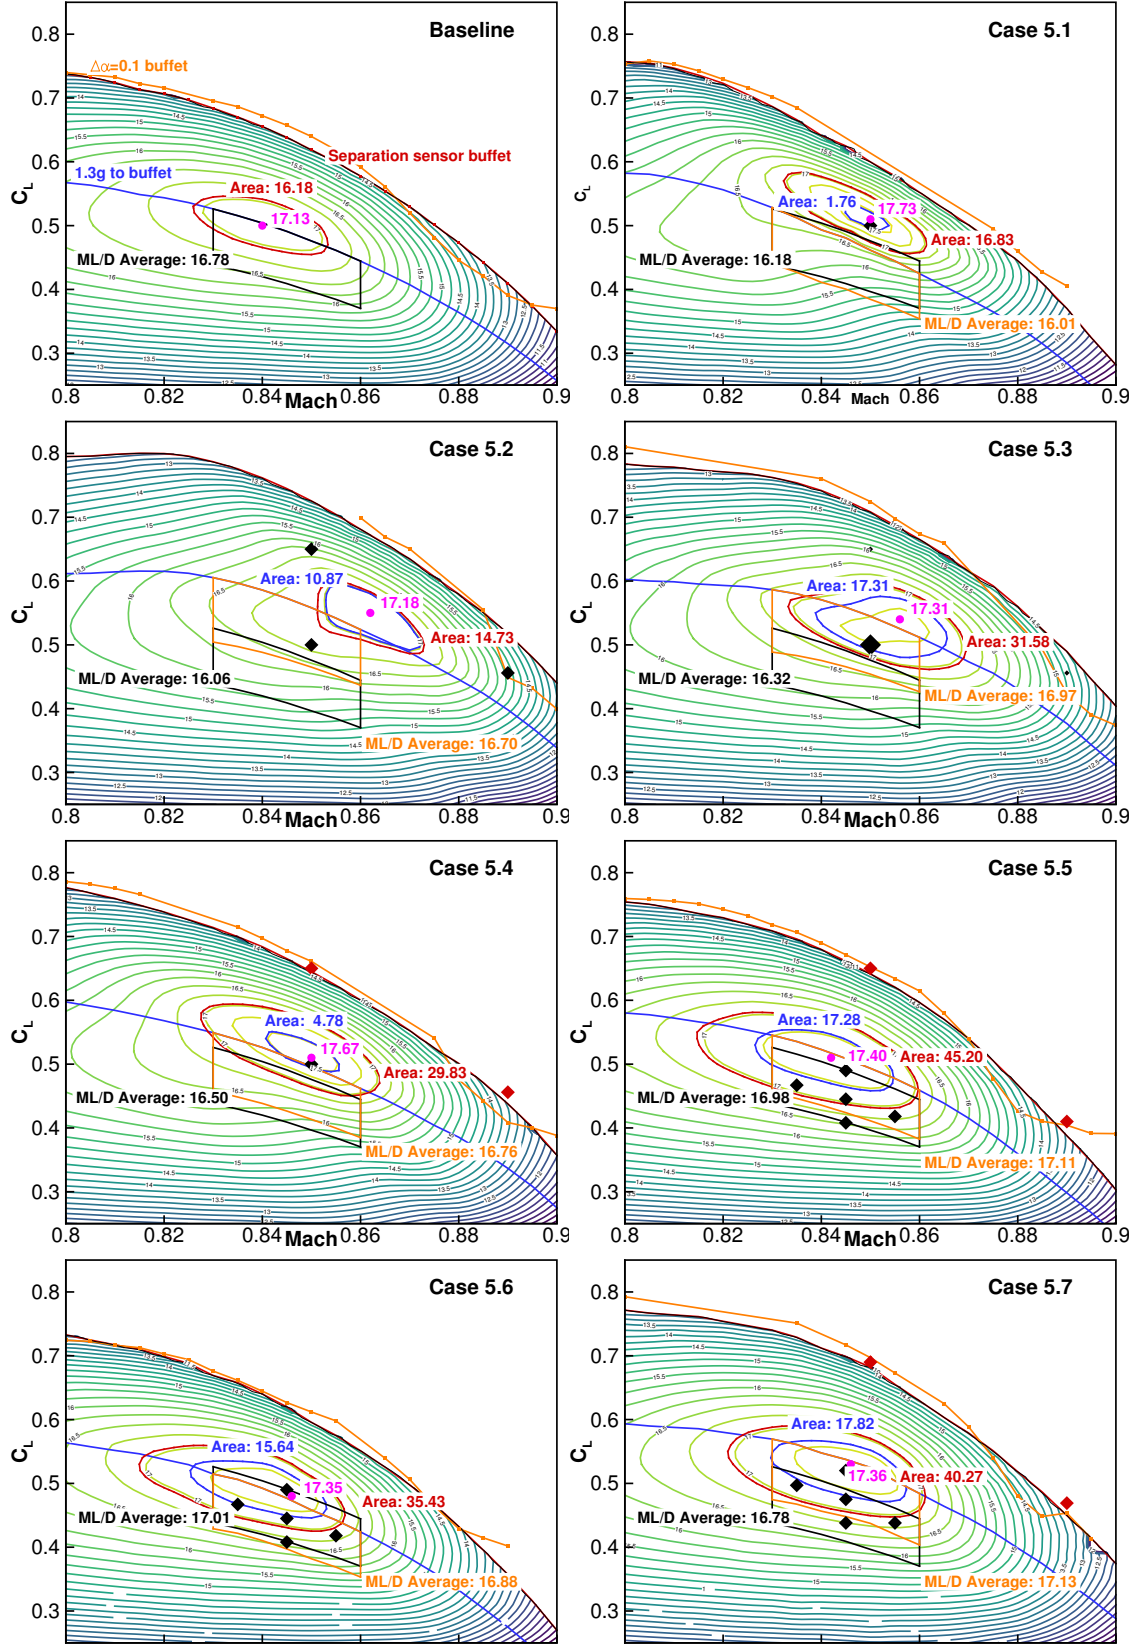
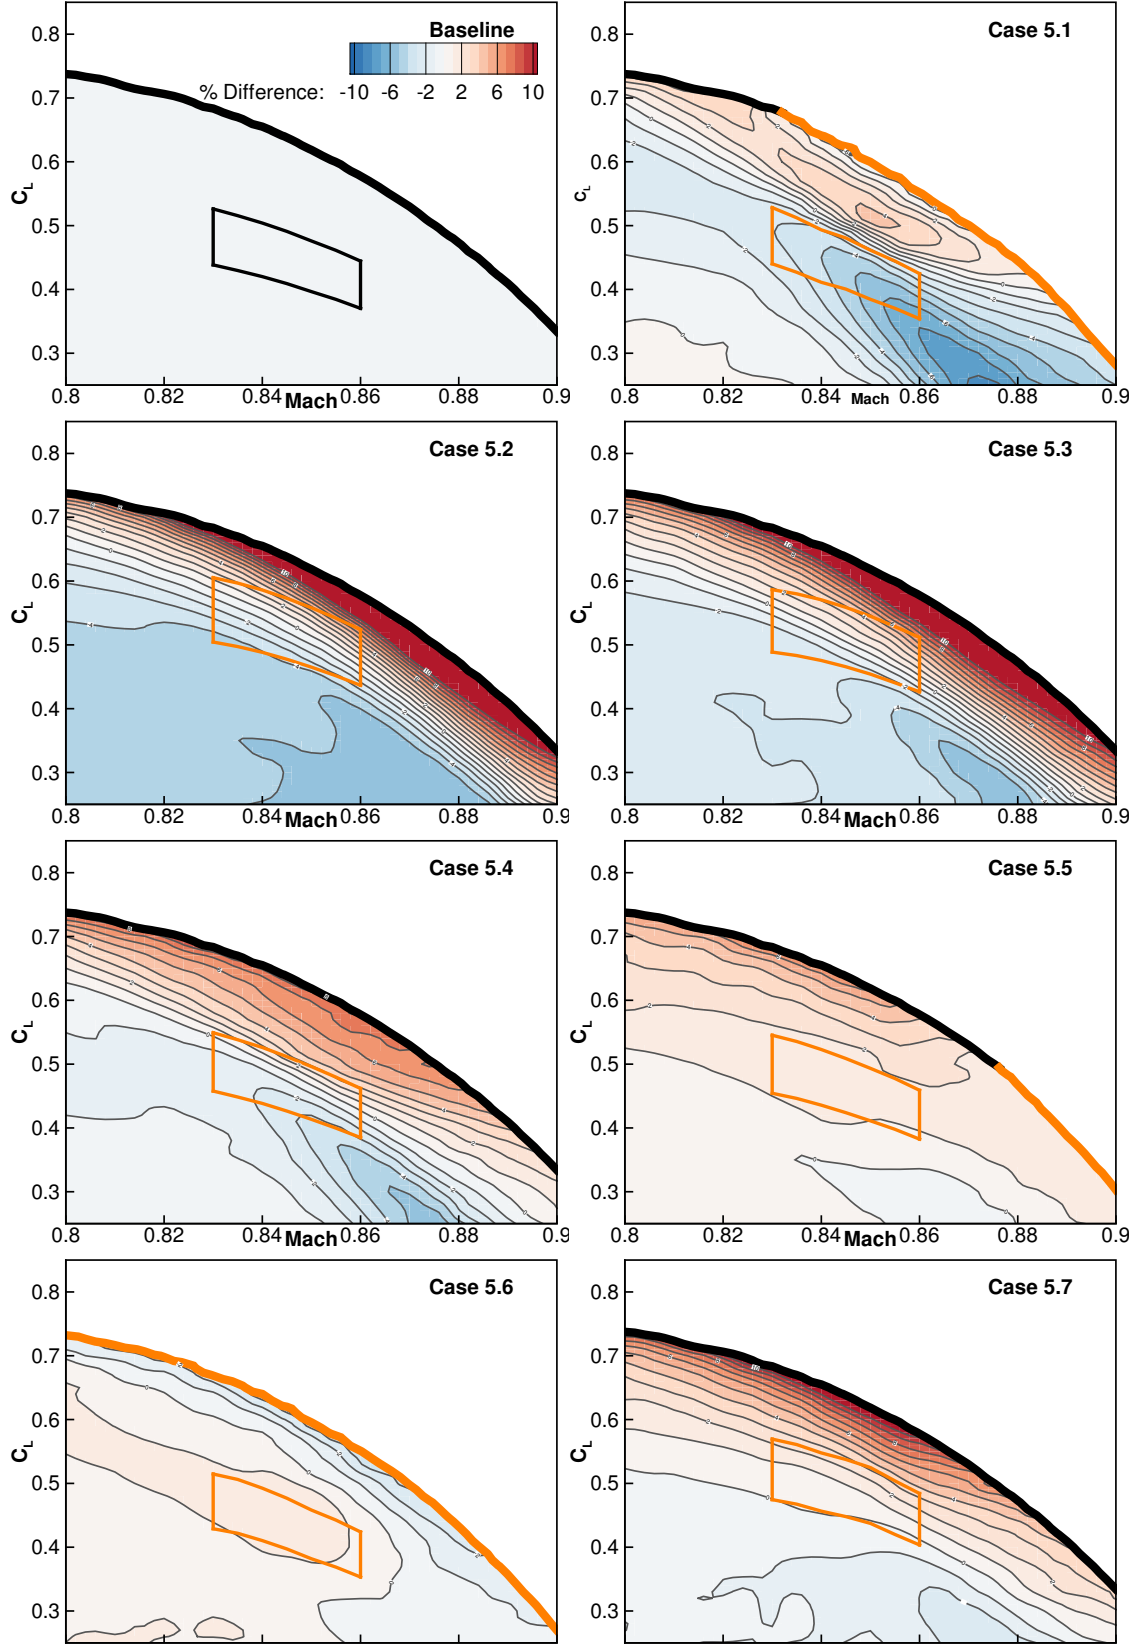Figure 16: Contours of ML/D for the baseline and for each optimized configuration.

Figure 17: Percent difference in ML/D between the baseline design and the optimized configuration.
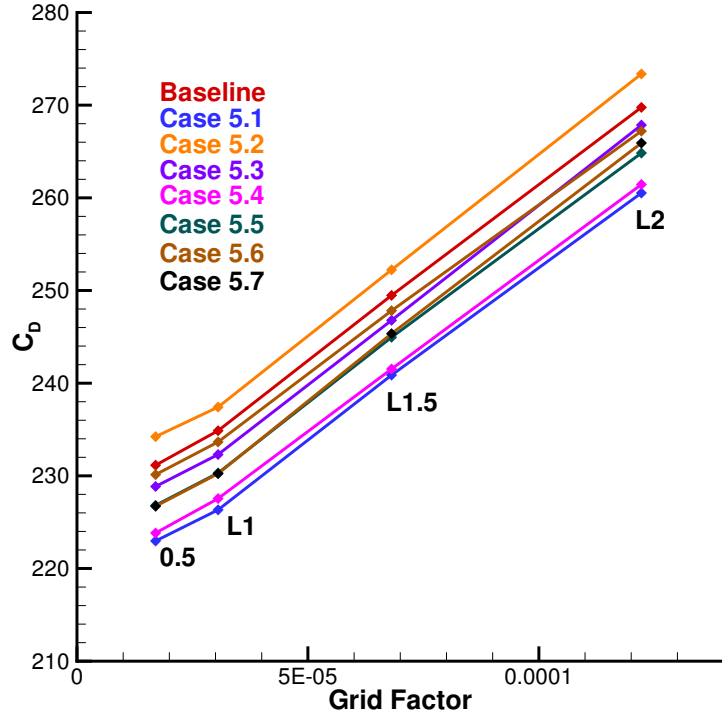
Figure 18: Grid convergence study for baseline and all optimized configurations, showing that the change in drag is constant between grid levels.

Table 4: Breakdown of computational cost in CPU hours.

Case	L2 optimization	L1.5 optimization	Contour	Grid convergence	Total
Baseline	–	–	1 346	817	2 162
5.1	289	611	1 270	1 009	3 179
5.2	2 378	2 394	1 795	1 121	7 688
5.3	1 290	2 505	1 750	910	6 457
5.4	1 507	2 602	1 384	1 024	6 518
5.5	2 090	3 506	1 392	830	7 369
5.6	1 111	1 803	1 147	610	4 673
5.7	4 136	6 623	1 800	696	13 255
Total	12 802	19 567	11 886	7 019	51 303

27

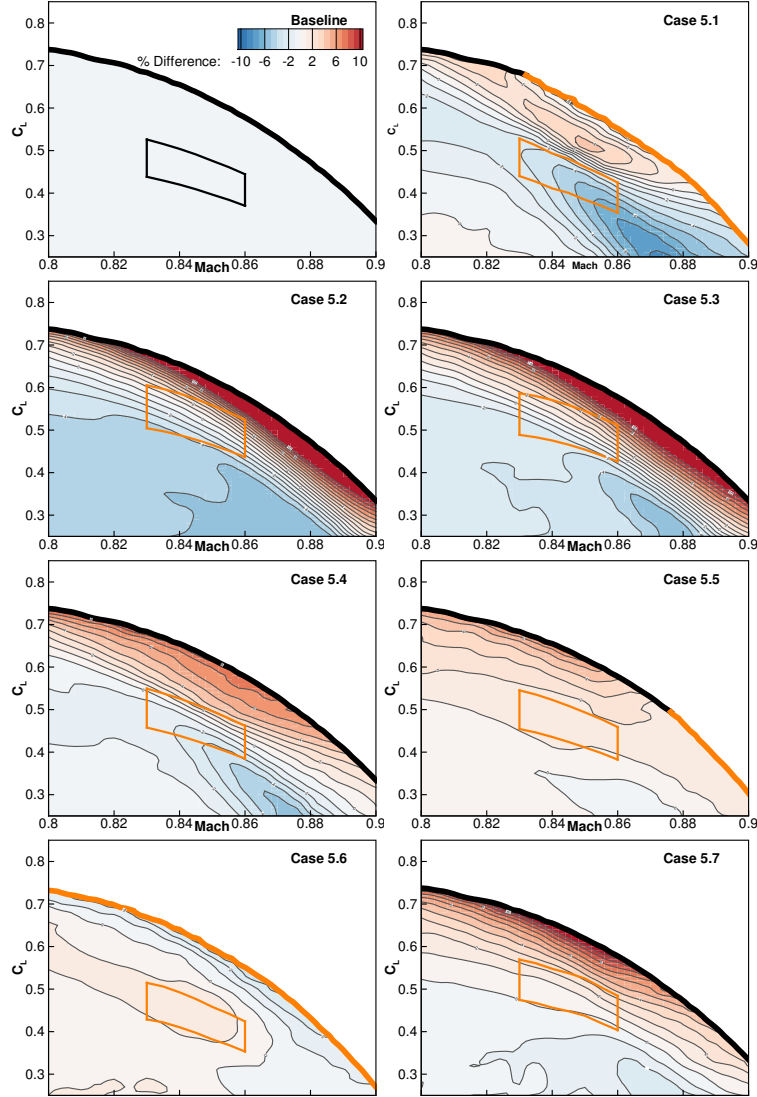Figure 19: Contours of the change in drag for the optimal designs relative to the baseline design show that this change is roughly constant for all designs between grid levels.

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] Hicks, R. M., Murman, E. M., and Vanderplaats, G. N., "An assessment of airfoil design by numerical optimization," Tech. rep., NASA TM X-3092, July 1974.

[4] Pironneau, O., "On Optimum Profiles in Stokes Flow," *Journal of Fluid Mechanics*, Vol. 59, No. 01, 1973, pp. 117–128. doi:10.1017/S002211207300145X.

[5] Jameson, A., "Aerodynamic Design via Control Theory," *Journal of Scientific Computing*, Vol. 3, No. 3, Sept. 1988, pp. 233–260.

[6] Peter, J. E. V. and Dwight, R. P., "Numerical Sensitivity Analysis for Aerodynamic Optimization: A Survey of Approaches," *Computers and Fluids*, Vol. 39, No. 3, March 2010, pp. 373–391. doi:10.1016/j.compfluid.2009.09.013.

[7] Nielsen, E. J. and Anderson, W. K., "Aerodynamic Design Optimization on Unstructured Meshes Using the Navier–Stokes Equations," *AIAA Journal*, Vol. 37, No. 11, 1999, pp. 1411–1419.

[8] Dwight, R. P. and Brezillon, J., "Efficient and robust algorithms for solution of the adjoint compressible Navier–Stokes equations with applications," *International Journal for Numerical Methods in Fluids*, Vol. 60, 2009, pp. 365–389. doi:10.1002/fld.1894.

[9] Hicken, J. E. and Zingg, D. W., "Aerodynamic Optimization Algorithm with Integrated Geometry Parameterization and Mesh Movement," *AIAA Journal*, Vol. 48, No. 2, Feb. 2009, pp. 400–413.

[10] 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.

[11] Lyu, Z. and Martins, J. R. R. A., "Aerodynamic Design Optimization Studies of a Blended-Wing-Body Aircraft," *Journal of Aircraft*, Vol. 51, No. 5, September 2014, pp. 1604–1617. doi:10.2514/1.C032491.

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

[13] Dumont, A. and Meheut, M., "Gradient-Based Optimization of CRM Wing-alone and Wing-body-tail Configurations by RANS Adjoint Technique," *54th AIAA Aerospace Sciences Meeting*, American Institute of Aeronautics and Astronautics (AIAA), January 2016. doi:10.2514/6.2016-1293.

[14] Kenway, G. K. W. and Martins, J. R. R. A., "Multipoint Aerodynamic Shape Optimization Investigations of the Common Research Model Wing," *AIAA Journal*, Vol. 54, No. 1, January 2016, pp. 113–128. doi:10.2514/1.J054154.

[15] Reuther, J. J., Jameson, A., Alonso, J. J., Rimlinger, M. J., and Saunders, D., "Constrained Multipoint Aerodynamic Shape Optimization Using an Adjoint Formulation and Parallel Computers, Part 1," *Journal of Aircraft*, Vol. 36, No. 1, 1999, pp. 51–60.

[16] Reuther, J. J., Jameson, A., Alonso, J. J., Rimlinger, M. J., and Saunders, D., "Constrained Multipoint Aerodynamic Shape Optimization Using an Adjoint Formulation and Parallel Computers, Part 2," *Journal of Aircraft*, Vol. 36, No. 1, 1999, pp. 61–74.

[17] Chen, S., Lyu, Z., Kenway, G. K. W., and Martins, J. R. R. A., "Aerodynamic Shape Optimization of the Common Research Model Wing-Body-Tail Configuration," *Journal of Aircraft*, Vol. 53, No. 1, January 2016, pp. 276–293. doi:10.2514/1.C033328.

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

[19] 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.

[20] 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.

[21] Kenway, G. K. W. and Martins, J. R. R. A., "AIAA ADODG Case 5: CRM Wing-Body-Tail Optimization at Flight Reynolds Number," Tech. rep., AIAA, May 2015.

[22] Ghazlane, I., Carrier, G., Dumont, A., and Désidéri, J.-A., "Aerostructural adjoint method for flexible wing optimization," *53rd AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference*, Honolulu, HI, April 2012, AIAA 2012-1924.

[23] 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. doi:10.2514/6.2014-3274, AIAA 2014-3274.

[24] Martins, J. R. R. A. and Lambe, A. B., "Multidisciplinary Design Optimization: A Survey of Architectures," *AIAA Journal*, Vol. 51, No. 9, September 2013, pp. 2049–2075. doi:10.2514/1.J051895.

[25] Martins, J. R. R. A. and Hwang, J. T., "Review and Unification of Methods for Computing Derivatives of Multidisciplinary Computational Models," *AIAA Journal*, Vol. 51, No. 11, November 2013, pp. 2582–2599. doi:10.2514/1.J052184.

[26] Wakayama, S., Page, M., and Liebeck, R., *Multidisciplinary optimization on an advanced composite wing*, American Institute of Aeronautics and Astronautics, 1996. doi:doi:10.2514/6.1996-4003.

[27] Wakayama, S., *Multidisciplinary design optimization of the blended-wing-body*, American Institute of Aeronautics and Astronautics, 2016/02/21 1998. doi:doi:10.2514/6.1998-4938.

[28] Bérard, A. and Isikveren, A. T., "Conceptual Design Prediction of the Buffet Envelope of Transport Aircraft," *Journal of Aircraft*, Vol. 46, No. 5, 2016/02/27 2009, pp. 1593–1606. doi:10.2514/1.41367.

[29] Vassberg, J. C. and Jameson, A., "Aerodynamic shape optimisation of a Reno race plane," *International Journal of Vehicle Design*, Vol. 28, No. 4, 2002, pp. 318–338.

[30] Jameson, A. and Vassberg, J., *Computational fluid dynamics for aerodynamic design - Its current and future impact*, American Institute of Aeronautics and Astronautics, 2016/02/27 2001. doi:doi:10.2514/6.2001-538.

[31] Buckley, H. P., Zhou, B. Y., and Zingg, D. W., "Airfoil Optimization Using Practical Aerodynamic Design Requirements," *Journal of Aircraft*, Vol. 47, No. 5, September 2010, pp. 1707–1719. doi:10.2514/1.C000256.

[32] Deck, S., "Numerical Simulation of Transonic Buffet over a Supercritical Airfoil," *AIAA Journal*, Vol. 43, No. 7, jul 2005, pp. 1556–1566. doi:10.2514/1.9885.

[33] Thiery, M. and Coustols, E., "URANS Computations of Shock-Induced Oscillations Over 2D Rigid Airfoils: Influence of Test Section Geometry," *Flow, Turbulence and Combustion*, Vol. 74, No. 4, 2005, pp. 331–354. doi:10.1007/s10494-005-0557-z.

[34] Crouch, J., Garbaruk, A., Magidov, D., and Travin, A., "Origin of transonic buffet on aerofoils," *Journal of fluid mechanics*, Vol. 628, 2009, pp. 357–369.

[35] Garnier, E. and Deck, S., *Direct and Large-Eddy Simulation VII: Proceedings of the Seventh International ERCOFTAC Workshop on Direct and Large-Eddy Simulation, held at the University of Trieste, September 8-10, 2008*, chap. Large-Eddy Simulation of Transonic Buffet over a Supercritical Airfoil, Springer Netherlands, Dordrecht, 2010, pp. 549–554. doi:10.1007/978-90-481-3652-0_81.

[36] Raveh, D. and Dowell, E., "Frequency lock-in phenomenon for oscillating airfoils in buffeting flows," *Journal of Fluids and Structures*, Vol. 27, No. 1, 2011, pp. 89–104. doi:http://dx.doi.org/10.1016/j.jfluidstructs.2010.10.001.

[37] Thomas, J. P. and Dowell, E. H., "Airfoil Transonic Flow Buffet Calculations Using the OVERFLOW 2 Flow Solver," *52nd AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics and Materials Conference*, April 2011, AIAA 2011-2077.

[38] Iovnovich, M. and Raveh, D. E., "Reynolds-Averaged Navier–Stokes Study of the Shock-Buffet Instability Mechanism," *AIAA Journal*, Vol. 50, No. 4, apr 2012, pp. 880–890. doi:10.2514/1.j051329.

[39] Thomas, J. and Dowell, E., *Discrete Adjoint Design Optimization Approach for Increasing Transonic Buffet Onset Angle-of-Attack*, American Institute of Aeronautics and Astronautics, 2016/02/21 2015. doi:doi:10.2514/6.2015-3435.

[40] Hall, K. C., Ekici, K., Thomas, J. P., and Dowell, E. H., "Harmonic balance methods applied to computational fluid dynamics problems," *International Journal of Computational Fluid Dynamics*, Vol. 27, No. 2, feb 2013, pp. 52–67. doi:10.1080/10618562.2012.742512.

[41] Mader, C. A. and Martins, J. R. R. A., "Derivatives for Time-Spectral Computational Fluid Dynamics Using an Automatic Differentiation Adjoint," *AIAA Journal*, Vol. 50, No. 12, December 2012, pp. 2809–2819. doi:10.2514/1.J051658.

[42] Lyu, Z. and Martins, J. R. R. A., "Aerodynamic Shape Optimization of an Adaptive Morphing Trailing Edge Wing," *Journal of Aircraft*, Vol. 52, No. 6, November 2015, pp. 1951–1970. doi:10.2514/1.C033116.

[43] Liem, R., Kenway, G. K. W., and Martins, J. R. R. A., "Multimission Aircraft Fuel Burn Minimization via Multipoint Aerostructural Optimization," *AIAA Journal*, Vol. 53, No. 1, January 2015, pp. 104–122. doi:10.2514/1.J052940.

[44] Brooks, T. R., Kennedy, G. J., and Martins, J. R. R. A., "High-fidelity Aerostructural Optimization of a High Aspect Ratio Tow-steered Wing," *57th AIAA/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference*, American Institute of Aeronautics and Astronautics, January 2016. doi:10.2514/6.2016-1179.

[45] Burdette, D. A., Kenway, G. K., and Martins, J. R. R. A., "Performance Evaluation of a Morphing Trailing Edge Using Multipoint Aerostructural Design Optimization," *57th AIAA/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference*, American Institute of Aeronautics and Astronautics, January 2016. doi:10.2514/6.2016-0159.

[46] Garg, N., Kenway, G. K. W., Lyu, Z., Martins, J. R. R. A., and Young, Y. L., "High-fidelity Hydrodynamic Shape Optimization of a 3-D Hydrofoil," *Journal of Ship Research*, Vol. 59, No. 4, December 2015, pp. 209–226. doi:10.5957/JOSR.59.4.150046.

[47] 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.

[48] 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.

[49] 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.

[50] Swanson, R. C., Radespiel, R., and Turkel, E., "Comparison of several dissipation algorithms for central difference schemes," *in AIAA 13th Computational Fluid Dynamics Conference, AIAA, 1997 (AIAA Paper*, pp. 97–1945.

[51] 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.

[52] Sederberg, T. W. and Parry, S. R., "Free-form Deformation of Solid Geometric Models," *SIGGRAPH Comput. Graph.*, Vol. 20, No. 4, Aug. 1986, pp. 151–160. doi:10.1145/15886.15903.

[53] Kenway, G. K., Kennedy, G. J., and Martins, J. R. R. A., "A CAD-Free Approach to High-Fidelity Aerostructural Optimization," *Proceedings of the 13th AIAA/ISSMO Multidisciplinary Analysis Optimization Conference*, Fort Worth, TX, Sept. 2010. doi:10.2514/6.2010-9231, AIAA 2010-9231.

[54] Mader, C. A. and Martins, J. R. R. A., "Computing Stability Derivatives and their Gradients for Aerodynamic Shape Optimization," *AIAA Journal*, Vol. 52, No. 11, November 2014, pp. 2533–2546. doi:10.2514/1.J052922.

[55] Kenway, G. K. W. and Martins, J. R. R. A., "Multipoint Aerodynamic Shape Optimization Investigations of the Common Research Model Wing," *Proceedings of the AIAA Science and Technology Forum and Exposition (SciTech)*, Kissimmee, FL, January 2015. doi:10.2514/6.2015-0264.

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

[57] 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.

[58] 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.

[59] Vos, R. and Farokhi, S., *Introduction to Transonic Aerodynamics*, Springer, 2015. doi:10.1007/978-94-017-9747-4.

[60] Rumsey, C. L., Allison, D. O., Biedron, R. T., Buning, P. G., Gainer, T. G., Morrison, J. H., Rivers, S. M., Mysko, S. J., and Witkowski, D. P., "CFD sensitivity analysis of a modern civil transport near buffet-onset conditions," Tech. Rep. TM-2001-211263, NASA, 2001.

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

[62] van Eijndhoven, J., *Buffet envelope prediction of transport aircraft during the conceptual design phase*, Master's thesis, Delft University of Technology, January 2012.

[63] Clark, R. and Pelkman, R., "High Reynolds number testing of advanced transport aircraft wings in the National Transonic Facility," *39th Aerospace Sciences Meeting and Exhibit*, American Institute of Aeronautics and Astronautics (AIAA), January 2001. doi:10.2514/6.2001-910.

[64] 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.

[65] 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.

[66] Garg, N., Lyu, Z., Dhert, T., Martins, J. R. R. A., and Young, Y. L., "High-fidelity Hydrodynamic Shape Optimization of a 3-D Morphing Hydrofoil," *Fourth International Symposium on Marine Propulsors*, June 2015.

[67] Garg, N., Kenway, G. K. W., Martins, J. R. R. A., and Young, Y. L., "High-fidelity Coupled Hydrostructural Optimization of a 3-D Hydrofoil," *International Symposium on Transport Phenomena and Dynamics of Rotating Machinery*, April 2016.

[68] Balakrishna, S. and Acheson, M., "Analysis of NASA Common Research Model Dynamic Data," *49th Aerospace Sciences Meeting Including the New Horizons Forum and Aerospace Exposition*, 2011. doi:10.2514/6.2011-1127.

[69] Vassberg, J. C., Tinoco, E. N., Mani, M., Rider, B., Zickuhr, T., Levy, D. W., Brodersen, O. P., Eisfeld, B., Crippa, S., Wahls, R. A., Morrison, J. H., Mavriplis, D. J., and Murayama, M., "Summary of the Fourth AIAA Computational Fluid Dynamics Drag Prediction Workshop," *Journal of Aircraft*, Vol. 51, No. 4, jul 2014, pp. 1070–1089. doi:10.2514/1.c032418.

[70] Haimes, R., "Automated Feature Extraction from Transient CFD Simulations," *Proceeding of the 7th Annual Conference of the CFD Society of Canada, Halifax, NS*, May 1999.

[71] Jameson, A., Schmidt, W., and Turkel, E., "Numerical Solution of the Euler Equations by Finite Volume Methods Using Runge-Kutta Time Stepping Schemes," *AIAA Paper* 81-1259, 1981.

[72] Towns, J., Cockerill, T., Dahan, M., Foster, I., Gaither, K., Grimshaw, A., Hazlewood, V., Lathrop, S., Lifka, D., Peterson, G. D., Roskies, R., Scott, J. R., and Wilkins-Diehr, N., "XSEDE: Accelerating Scientific Discovery," *Computing in Science & Engineering*, Vol. 16, No. 5, September 2014, pp. 62–74. doi:10.1109/MCSE.2014.80.

[73] Akima, H., "A new method of interpolation and smooth curve fitting based on local procedures," *Journal of the ACM*, Vol. 4, No. 17, 1970, pp. 589–602.

A Generation of Contour Plots

The generation of the contour plots shown in Fig. 16 warrants further explanation. These contours are not simple α sweeps, because the tail angle that gives $C_M = 0$ must be determined at each point. Once the flight condition in the contour is trimmed, the trim drag penalty is included in the computed drag coefficient. Naively performing a secant search to determine the tail angle at each point would require at least three CFD solutions. However, because each contour plot requires approximately 400 trim-converged solutions, we seek instead an alternative approach to reduce the computational cost to the extent possible.

One way to reduce the computational cost of producing these contours is to reuse previously evaluated points to continually update the 2×2 Jacobian of the residual, $F = [C_L - C_L^*, C_M]$, with respect to $[\alpha, \eta]$, where α is the angle of attack and η is the tail rotation angle. With an accurate Jacobian, we can use Newton's method to simultaneously determine the new α and η required to produce a trimmed solution at a new C_L. The full procedure is listed in Algorithm 2. An auxiliary function for computing the residual for a given (α, η) is given in Algorithm 1.

In practice, only one subiteration is necessary for most points, because the C_L and C_M functions are not highly nonlinear functions of α and η over most of the contour region. Generally, additional subiterations are only necessary as buffet is approached due to the more-rapid variation in the lift curve slope. For example, the contour for the baseline configuration requires 430 function evaluations to produce 356 converged trimmed-C_L solutions, an increase of only 20%. Note that we only check for the convergence of C_M, because precisely matching the lift coefficients to the specified target is not critical.

Once all the raw data are generated, C_L, C_D, C_M, and the separation sensor values are interpolated by using an Akima spline [73] to produce a regular M-C_L grid. This regular grid is then used for further computations, such as the difference plots shown in Fig. 17 and the drag-divergence curves in Figs. 13 and 14, and for extracting particular contours and computing the average performance over specific integration regions.

Algorithm 1 Trim-C_L function

1: **function** $F(x, C_L^*)$
2:     Set $\alpha = x[0]$ and $\eta = x[1]$
3:     Solve CFD problem
4: **return** $F = [C_L - C_L^*, C_M]^T$
5: **end function**

Algorithm 2 Algorithm to compute trimmed-C_L contours

1: Given: $C_{L_{\min}}$, $[M_1, ..., M_N]$, ΔC_L^*, α_0, η_0, $\Delta\alpha$, $\Delta\eta$, Sep_{max}, maxSubIter, tol
2: Set $M = M_1$
3: $F_0 = F([\alpha_0, \eta_0]^T, C_{L_{\min}})$
4: $F_1 = F([\alpha_0 + \Delta\alpha, \eta_0]^T, C_{L_{\min}})$
5: $F_2 = F([\alpha_0, \eta_0 + \Delta\eta]^T, C_{L_{\min}})$
6: $J = \left[\frac{F_1 - F_0}{\Delta\alpha} \ \frac{F_2 - F_0}{\Delta\eta} \right]$ ▷ Initial finite-difference Jacobian approximation
7: $J_{\text{save}} = J$
8: $x_{\text{save}} = [\alpha_0, \eta_0]^T$
9: **for** $k \leftarrow 1, N$ **do**   ▷ Loop over N sequential Mach numbers
10:     set $M = M_k$
11:     $J = J_{\text{save}}$ ▷ Restore Jacobian for lowest C_L
12:     $x_n = x_{\text{save}}$ ▷ Restore x for lowest C_L
13:     $\Delta C_L = \Delta C_L^*$ ▷ Restore target C_L increment
14:     $C_L^* = C_{L_{\min}}$ ▷ Set the target C_L to the lowest desired C_L
15:     $F_n = F(x_n, C_L^*)$ ▷ Evaluate current point
16:     continue = True
17:     $i_\alpha = 0$
18:     **while** continue **do**   ▷ α increment loop
19:        **for** $j \leftarrow 1, \text{maxSubIter}$ **do**   ▷ Subiteration loop
20:          $dx = J^{-1} F_n$ ▷ Newton's method for update
21:          $x_{n+1} = x_n - dx$ ▷ New (α, η) solution
22:          $F_{n+1} = F(x_{n+1}, C_L^*)$ ▷ Solve CFD for the new x
23:          $dF = F_{n+1} - F_n$
24:          $J = J + \frac{dF - J dx}{\|dx\|^2} dx^T$ ▷ Update Jacobian using Broyden's method
25:          $x_n = x_{n+1}$ ▷ Set x_n for next subiteration
26:          **if** $i_\alpha = 0$ **then**   ▷ Store J and x_{n+1} for the next Mach number on the lowest C_L
27:             $J_{\text{save}} = J$
28:             $x_{\text{save}} = x_n$
29:          **end if**
30:          Evaluate separation sensor, Sep
31:          **if** Sep > 0.01 **then**
32:             $\Delta C_L = 0.01$ ▷ Reduce the target ΔC_L as buffet is approached
33:          **end if**
34:          **if** $\text{abs}(F_n[1]) < \text{tol}$ or $j = \text{maxSubIter}$ **then**   ▷ Only check convergence for moment
35:             $C_L^* = C_L^* + \Delta C_L$ ▷ Update the next target C_L
36:             $F_n = [F_{n+1}[0] - C_L^*, F_{n+1}[1]]$ ▷ Set the function value for the next C_L
37:             Break subiteration loop
38:          **else**
39:             $F_n = F_{n+1}$ ▷ Continue to refine the current target C_L
40:          **end if**
41:        **end for**
42:        $i_\alpha = i_\alpha + 1$
43:        **if** Sep > $\text{Sep}_{\max}$ **then**
44:          continue=False   ▷ Reached buffet onset, so proceed to next Mach number
45:        **end if**
46:     **end while**
47: **end for**