Multidisciplinary Design Optimization of Aircraft Configurations
Part 2: High-fidelity aerostructural optimization

Joaquim R. R. A. Martins, Gaetan K. W. Kenway, and Timothy Brooks
University of Michigan

Contents

1 Introduction

When it comes to flight, the wing is arguably the most crucial component. As the legendary Boeing aircraft designer Jack Steiner put it, "The wing is where you're going to fail.". In a book detailing the origins of the Boeing 747 Jumbo Jet, Irving (1993) writes:

> "Designing the wing involved literally thousands of decisions that could add up to an invaluable asset, a proprietary store of knowledge. A competitor could look at the wing, measure it even, and make a good guess about its internal structure. But a wing has as many invisible tricks built into its shape as a Savile Row suit; you would need to tear it apart and study every strand to figure out its secrets."

These "invisible tricks" are a reflection of the complexity involved wing modeling and design, where flexibility couples the aerodynamic shape to the structural layout and sizing. Thus it is of the utmost importance to consider the coupling between aerodynamic and structural models, not only in analysis, but also in design. The "thousands of decisions" cited in the above quote can be mapped to design variables, which involve both aerodynamic shape, and structural design variables. The aerodynamic flow is sensitive to the slightest change in aerodynamic shape, so it is important to parameterize the aerodynamic shape and structural sizing with a large number variables.

When optimizing both aerodynamics and structures, we need to consider the effect of the aerodynamic shape variables and structural sizing variables on the weight, which also affects the fuel burn. Thus, there are complex multidisciplinary trade-offs. Numerical optimization is a powerful tool that can perform these trade-offs automatically. Aerospace engineering researchers recognized this as soon as multi-physics models for wings were available, establishing the field of multidisciplinary design optimization (MDO) (Haftka, 1977; Martins and Lambe, 2013).

In Part 1 of this lecture we introduced a general framework for the solution of multidisciplinary system and assembling the coupled derivative using an adjoint method (MAUD). Here, we cover the MDO of aircraft configurations based on high-fidelity models in more detail, with emphasis on the coupling of the aerodynamics and structures disciplines. We show a wing design example where we tackle the compounding challenges of modeling the wing with high fidelity, while optimizing it with respect to hundreds of design variables. We are able to do this successfully through the use of high-performance parallel computing, fast solvers, state-of-the-art gradient-based optimization, and an efficient approach for computing the coupled derivatives for the aerostructural solver.

We start by reviewing our computational fluid dynamics (CFD) solver and explain how the adjoint for this solver was developed and implemented in Sec. 2. We then review two aerodynamic shape optimization problems whose solution was enabled by this solver in Sec. 3. In Sec. 4 we expand the problem by coupling the aerodynamics to a structural solver, and develop the corresponding coupled adjoint equations. Several applications of this aerostructural analysis and design optimization framework follow in Sec. 5.

2 Aerodynamic shape optimization framework

The aerodynamic shape optimization procedure is summarized in Fig. 1, which shows each of the modules described in this section. The optimization loop relies on a gradient-based

optimization algorithm to decide on the new set of design variables x at each iteration. For each new set of design variables, we must generate a new geometry and corresponding CFD mesh, so that the flow solver can compute the drag, lift, and moment coefficients for the new shape. The adjoint solver then computes the derivatives required by the optimization algorithm to perform the next design optimization iteration. This framework has been successfully used to perform several studies (Chen et al., 2016; Kenway and Martins, 2015b; Lyu et al., 2015; Lyu and Martins, 2014; Mader and Martins, 2013, 2014), some of which we review in Sec. 3.
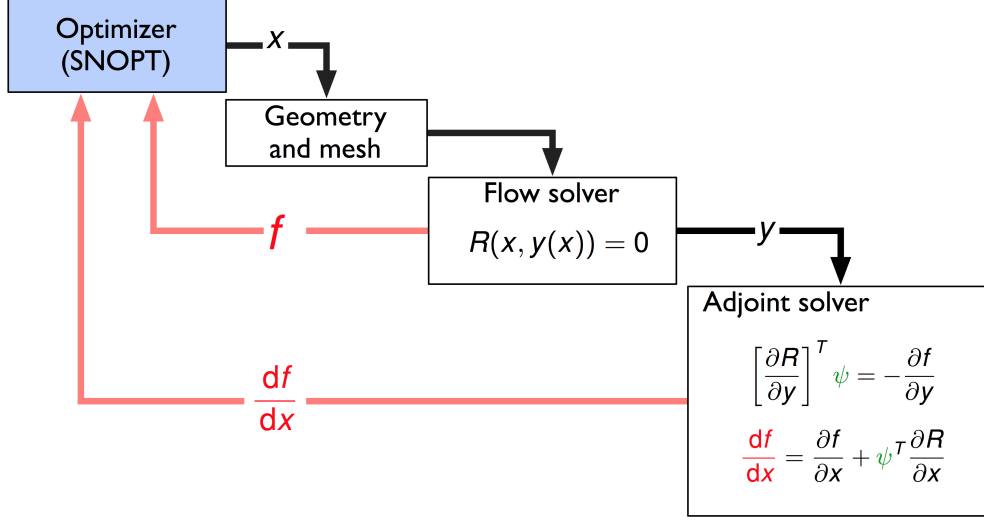

Figure 1: Aerodynamic shape optimization iteration loop.

2.1 CFD solver

The CFD solver used in the work herein is SUMad, which is a finite-volume, cell-centered multiblock solver for the compressible Euler, laminar Navier–Stokes, and RANS equations (steady, unsteady, and time-periodic modes) (van der Weide et al., 2006). The spatial discretization follows the Jameson–Schmidt–Turkel scheme (Jameson et al., 1981). Central difference augmented with artificial dissipation is used for the discretization of the inviscid fluxes and viscous fluxes use a central discretization. SUMad provides options for a variety of turbulence models with one, two, or four equations and options for adaptive wall functions. The Spalart–Allmaras model (Spalart and Allmaras, 1992) is used herein unless otherwise noted.

The solver was originally developed as a Runge–Kutta (RK) solver employing well-known steady-state acceleration techniques including local time-stepping, implicit residual smoothing, and selective evaluation of viscous fluxes on alternate stages. For RANS analysis, the turbulent equations are solved in loosely coupled fashion using a diagonally-dominant alternating direction implicit (DD-ADI) scheme with implicit boundary condition treatment. Acceptable load balancing on large processor counts is achieved using a built-in automatic block splitting algorithm.

Recently, two additional solution techniques were added to the solver: A diagonalized

ADI method for the mean-flow equations, and a fully coupled Newton–Krylov (NK) solver. Depending on the particular problem at hand, the DADI method is 3–5 times faster than the RK method and the NK method is 8–25 times faster.

In practice, we use the RK solver at the start of a solution and then switch to NK for the final convergence. This provides an excellent trade-off between robustness and computational efficiency.

The steady state solution of the RANS equations can be expressed as finding the flow state w such that the residuals of the discretized partial differential equations for each control volume are zero, i.e.,

$$R(x, w) = 0, \tag{1}$$

where x is the vector of design variables, which is fixed for each flow solution. Both R and w are vectors of dimension $(5 + n_{\text{turb}}) \times n_{\text{cells}}$, where n_{turb} is the number of turbulence model equations, and n_{cells} is the number of cells in the CFD mesh.

2.2 Geometric parametrization

In this work we use a free-form deformation (FFD) volume approach (Sederberg and Parry, 1986) that we implemented (Kenway et al., 2010) and have been using extensively for aerodynamic (Chen et al., 2016; Kenway and Martins, 2015b; Lyu et al., 2015; Lyu and Martins, 2014; Mader and Martins, 2013, 2014) and aerostructural optimization studies (Kenway et al., 2014a,b; Kenway and Martins, 2014; Liem et al., 2015).

The FFD approach can be visualized as embedding the spatial coordinates defining a geometry inside a flexible volume. The parametric locations corresponding to the baseline geometry are found using a Newton search algorithm. Once the baseline geometry has been embedded, perturbations made to the FFD volume propagate within the embedded geometry via the evaluation of the nodes at their parametric locations. An example of an FFD for a wing shape optimization is shown in Fig. 2.
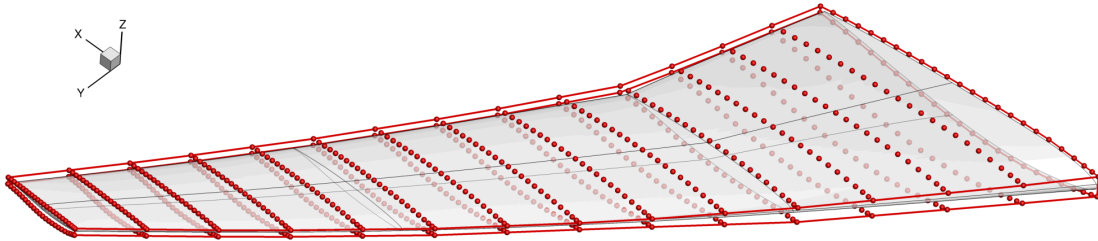
Figure 2: The shape design variables are FFD control point z-displacements (red spheres).

2.3 Mesh movement

The FFD approach used to parametrize the geometry applies deformations only to the surface mesh—the part of the volume mesh that lies on the physical surface. A mesh movement procedure is then required to propagate surface perturbations to the remainder of the CFD volume mesh. The mesh movement algorithm used in this work is an efficient analytic inverse distance method similar to the one described by Luke et al. (2012). Updating the mesh for a new configuration is fast, typically requiring less than 0.1% of the CFD solution time.

4

2.4 Automatic differentiation adjoint

As previously mentioned, accurate and efficient computation of derivatives of objective and constraint functions with respect to large numbers of design variables is crucial in the efficiency of the overall optimization.

In Part 1 of this lecture, we introduced the adjoint method, which computes derivatives of a function of interest at a cost independent of the number of design variables. The computational cost of the adjoint method scales with the number of functions of interest, which is usually much lower than the number of design variables for aerodynamic shape optimization problems.

As derived in Part 1, the discrete adjoint equations can be written as,

$$\left[\frac{\partial R}{\partial w}\right]^T \psi = -\left[\frac{\partial F}{\partial w}\right]^T, \tag{2}$$

where we replaced the general state vector y by the vector of flow state variables w. The function of interest F, in the case of aerodynamic shape optimization, is usually the drag coefficient as the objective. Additional adjoint equations must also be solved when F is a constraint function that depends on the flow solution—usually lift and moment coefficients. Once the adjoint vector ψ is computed, the total derivatives can be obtained using,

$$\frac{\mathrm{d}f}{\mathrm{d}x} = \frac{\partial F}{\partial x} + \psi^T \frac{\partial R}{\partial x}. \tag{3}$$

Again, as explained in Part 1, the partial derivatives in the equations represent an explicit dependence that do not require convergence of the residual (the flow solution in this case). For example, if x is a set of wing shape variables, $\partial R/\partial x$ represents the sensitivity of the residuals at each cell to changes in the mesh due to wing shape changes, where the flow solution w is kept fixed.

Traditionally, the differentiation for these partial derivatives is done by hand, and the code that computes them is generated manually, which requires a long development time and is prone to errors. In some cases, researchers have also used finite differences and the complex-step method. Our approach is to use automatic differentiation (AD) to produce the code that computes these partial derivatives. As explained in Part 1, AD is a tool that given the source code of a program, automatically generates code that computes the derivatives of specified outputs to specified inputs. In its pure form, we would use AD on the whole CFD code including the iteration procedure. Using this brute force approach blindly incurs a large computational cost both in terms of memory and time, although there have been efforts towards making this feasible (Albring et al., 2016). By selectively applying AD to compute the terms in Eqs. (2) and (3), we strike a balance between computational efficiency and development time.

This AD adjoint (ADjoint) approach was first demonstrated by Mader et al. (2008). Since this demonstration, we have developed different ADjoint approaches and refined the implementation to be both more automated, and more efficient. This idea was reinforced in the survey on CFD-based derivative computation by Peter and Dwight (2010), who wrote "*It is our belief that, at least in the medium term, industrially relevant linearized codes will be developed by using AD to differentiate individual nonlinear routines, which are then assembled by an expert.*"

In our experience, the desirable features for an ADjoint approach are:

Consistency: The code must yield derivatives consistent with the flow solution and be verifiable. Typically we verify it against the complex-step approach and hence are able to find small errors. Such errors cannot be found when verifying against finite-differences, and although small, the errors could become larger in different cases.

Maintainability: For ease of maintenance, the approach should not require modification or duplication of the original code.

Efficiency: The ADjoint code should be efficient both in terms of memory usage and computational time; it should also incur no penalty to the original CFD solution code.

Ease of implementation: The implementation should be as automatic as possible, with short development time.

In the last decade, we have experimented with different ADjoint approaches that address these feature to various degrees.

Single cell residual: This consists in using AD on a single cell residual routine, and loop over cells to assemble the full Jacobian (Mader et al., 2008; Martins et al., 2006).

Forward mode coloring: The complete residual routine is differentiated with AD using coloring for efficiency and the full Jacobian is stored. (Lyu et al., 2013).

Full reverse mode: A master routine is created such that when AD is applied to it, it yields the desired Jacobian-vector products and derivatives.

Hybrid reverse mode: AD is applied to individual subroutines in the master routine and the Jacobian-vector products are manually assembled using the chain rule.

The last two approaches are the most sophisticated and do the best in fulfilling the desirable features listed above.

The single cell approach was the first demonstration of the ADjoint approach, and although it was successfully used for aerodynamic shape optimization (Mader and Martins, 2013), it required code duplication, making it a poor approach from the maintainability point of view. The forward mode coloring approach enabled us to develop a more efficient adjoint that was much more maintainable because AD was applied to the original residual computation routine. Using coloring techniques, we were able to reduce the computational cost of forward mode AD (Lyu et al., 2013). However, this approach requires storing the full Jacobian $\partial R/\partial w$, which means that a direct solver can be used, but very large cases might be memory limited, depending on the computer memory available. This approach was used to solve several aerodynamic shape optimization problems (Lyu et al., 2015; Lyu and Martins, 2014, 2015).

The last two approaches have the advantage that they do not require storing the flow Jacobian $\partial R/\partial w$. Instead, they yield the transpose matrix-vector product of the Jacobian with a specified vector v, i.e., $[\partial R/\partial w]^T v$. This product is then used to efficiently solve

the adjoint equations using an iterative solver, such as the generalized minimum residual (GMRES) method.

The key to these two approaches is to create a routine that maps the design variables and flow variables to the CFD residuals, as well as the functions of interest, i.e., $(x, w) \mapsto (R(x, w), F(x, w))$. This mapping must not include iterations to solve the CFD equations, only explicit dependencies. This master routine is not used directly, but is created for the sole purpose of being differentiated with AD. The resulting AD code computes all the partial derivatives needed in the adjoint equation (2) and the total derivative computation (3) using reverse mode AD.

The adjoint linearization (2) includes both the mean flow equations as well as turbulence model. This results in more accurate derivatives for design optimization but yields adjoint systems that are ill-conditioned and thus more challenging to solve. We use the solvers in PETSc (Balay et al., 1997) to solve the adjoint equations. The solution procedure is based on a GMRES outer loop. By experimenting with the PETSc solution options, we found that using a restricted additive Schwartz preconditioner with an overlap of 2 and an ILU factorization with fill levels of 2 or 3 on each sub-domain results in the lowest overall CPU time.

Strong scaling performance of both the flow solution (using the NK solver) and adjoint solutions are shown in Fig. 2.4 for a coarse mesh (450k cells; typical for our exploratory optimization) and a finer one (3.6M cells; typical for refining the optimized results). Both
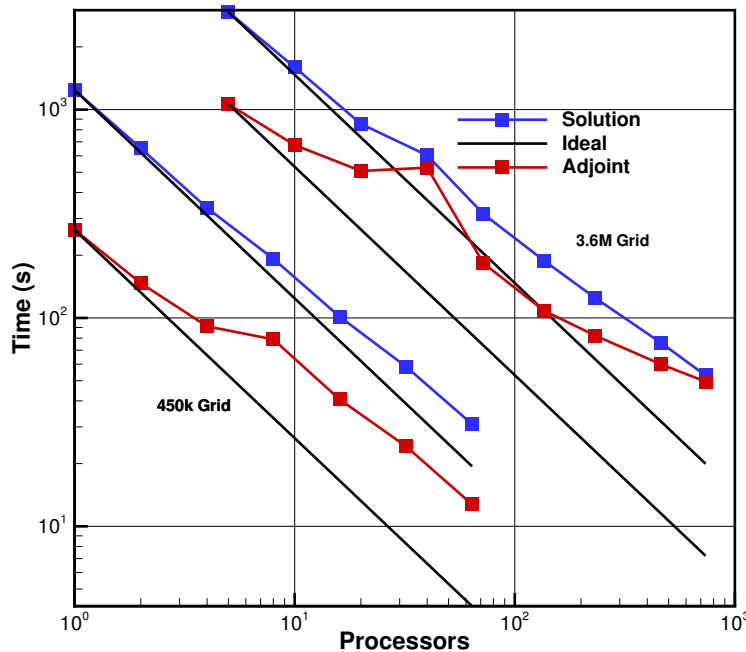

Figure 3: The RANS solver and adjoint solver exhibit strong parallel scaling.

meshes show good speedup for the NK flow solver through two orders of magnitude change in processor count. The adjoint solver shows good scaling, although not quite as good as the flow solver technique.

2.5 Optimization algorithm

The optimization algorithm we use for all the results presented herein is SNOPT (sparse nonlinear optimizer) (Gill et al., 2002) through the Python interface pyOpt (Perez et al., 2012). 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 a smooth augmented Lagrangian merit function, and the Hessian of the Lagrangian is approximated using a limited-memory quasi-Newton method.

3 Aerodynamic shape optimization applications

We now present a few of the results that we have obtained with the aerodynamic shape optimization framework detailed in the previous section. We start by summarizing the solution of a benchmark RANS-based wing shape optimization problem in Sec. 3.1, and then present results for a full aircraft configuration including trim and buffet constraints in Sec. 3.2.

3.1 Common Research Model wing

Until recently, despite considerable research on aerodynamic shape optimization, there had been no standard benchmark problems allowing researchers to compare results. The AIAA Aerodynamic Design Optimization Discussion Group (ADODG) defined several benchmark cases including a 3D RANS-based case. The baseline geometry for this case is a wing with a blunt trailing edge extracted from the CRM wing-body geometry (Vassberg, 2011; Vassberg et al., 2008).

The results presented herein are from Lyu et al. (2015), who have all geometries meshes publicly available. The aerodynamic shape optimization seeks to minimize the drag coefficient by varying the shape design variables subject to a lift constraint and a pitching moment constraint. The shape design variables are the z-coordinate movements of 720 control points on the FFD volume (shown in Fig. 2) and the angle-of-attack. There are 750 thickness constraints imposed in a 25 chordwise and 30 spanwise grid covering the full span and from 1% to 99% local chord. The thickness is set to be greater than 25% of the baseline thickness at each location. Finally, the internal volume is constrained to be greater than or equal to the baseline volume. The complete optimization problem is described in Table 3.1.

	Function/variable	Description	Quantity
minimize	C_D	Drag coefficient	
with respect to	α	Angle of attack	1
	z	FFD control point z-coordinates	720
		Total design variables	721
subject to	$C_L = 0.5$	Lift coefficient constraint	1
	$C_{M_y} \geq -0.17$	Moment coefficient constraint	1
	$t \geq 0.25 t_{\text{base}}$	Minimum thickness constraints	750
	$V \geq V_{\text{base}}$	Minimum volume constraint	1
	$\Delta z_{\text{TE,upper}} = -\Delta z_{\text{TE,lower}}$	Fixed trailing edge constraints	15
	$\Delta z_{\text{LE,upper,root}} = -\Delta z_{\text{LE,lower,root}}$	Fixed wing root incidence constraint	1
		Total constraints	769

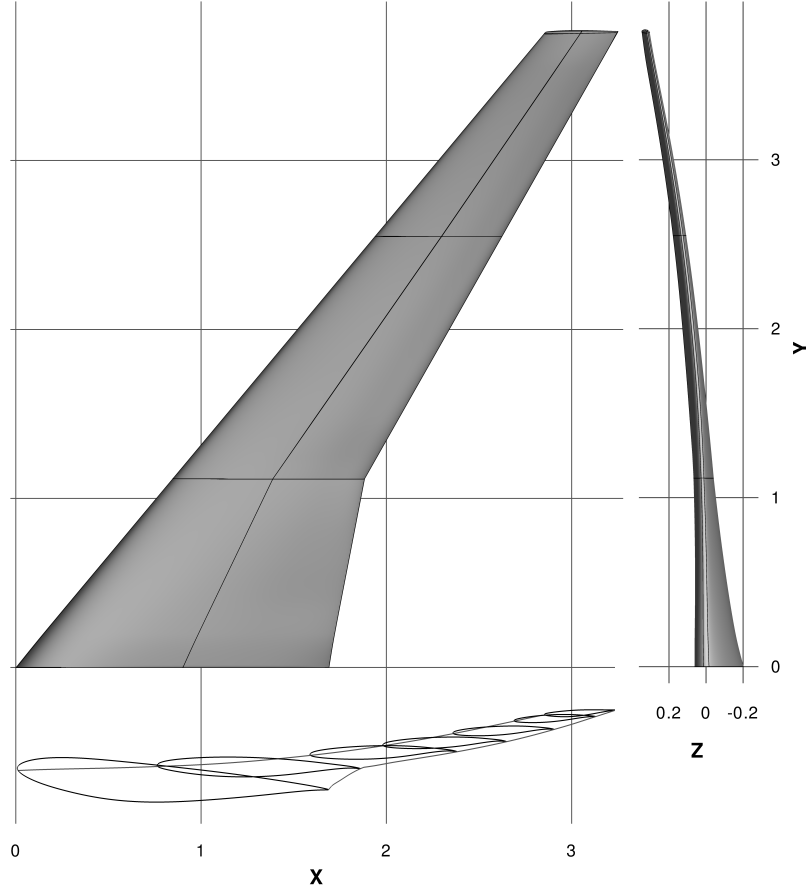Table 1: Aerodynamic shape optimization problem

Figure 4: Baseline CRM wing geometry scaled by its mean aerodynamic chord.

Figure 5 shows a detailed comparison of the baseline wing and the optimized wing. In this figure, the baseline wing results are shown in red and the optimized wing results are shown in blue. At the optimum, the lift coefficient target is met, and the pitching moment is reduced to the lowest allowed value. The lift distribution of the optimized wing is much closer to the elliptical distribution than that of the baseline, indicating an induced drag that is close to the theoretical minimum for a planar wake. This is achieved by fine-tuning the twist distribution and airfoil shapes.

The optimized thickness distribution is significantly different from that of the baseline, since the thicknesses are allowed to decrease to 25% of the original thickness, and there is a strong incentive to reduce the airfoil thicknesses in order to reduce wave drag. The volume is constrained to be greater than or equal to the baseline volume, so the optimizer drastically decreases the thickness of the airfoils on the outboard of the wing to the lower bounds, where there is less volume to be gained, while increasing the thickness near the root (up to 20%), where the chords are larger and the volume-drag trade-off is more favorable. The low outboard thickness would in practice incur a large structural weight penalty, and to trade off the reduction in drag and increase in weight would require aerostructural optimization, which we address in Sec. 5.

The baseline wing exhibits a front of closely spaced pressure contour lines spanning a sig-
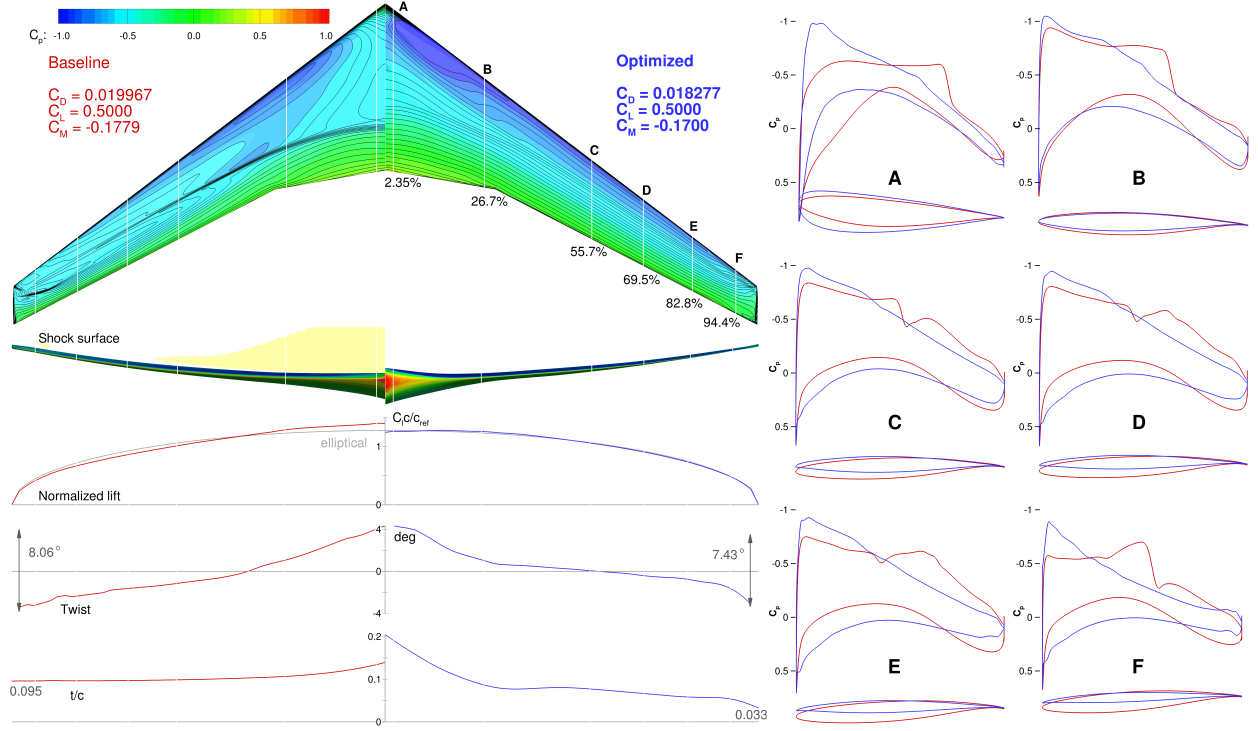
9

Figure 5: The optimized wing is shock-free and has 8.5% lower drag.

nificant portion of the wing, indicating a shock. The optimized wing shows parallel pressure contour lines with uniform spacing, indicating a shock-free solution under the nominal flight condition. This is confirmed by the shock surface plots: we can see that the baseline wing has a shock on the upper surface, while the optimized wing does not show shocks under the design condition. The shock elimination can also be seen on the airfoil C_p distributions. The sharp increase in local pressure due to the shock becomes a gradual change from the leading edge to the trailing edge.

To ensure that the result of our single-point optimization has sufficient accuracy, we conducted a grid convergence study of the optimized design. The mesh convergence plot for both the baseline and optimized geometry meshes is shown in Fig. 6. The zero-grid spacing drag, which was obtained using Richardson's extrapolation, is also plotted in the figure. We can see that the L0 mesh has sufficient accuracy: the difference in the drag coefficient for the L0 mesh and the value obtained for the zero-grid spacing is within one drag count. The variation in drag coefficient between the baseline and optimized meshes is nearly constant for each grid level, which gives us confidence that the optimization using the coarse meshes represent the design space trends sufficiently well. Therefore, we perform the remaining optimization studies on the coarser mesh (L2), assuming that we capture the correct design trends.

Using this benchmark, Lyu et al. (2015) determine that there is no significant multimodality by showing that the optimizers always converges to the same geometry even when starting from a wide variety of starting wing shapes. One such example is shown in Fig.7, where the starting geometry is a random perturbation of the baseline geometry, and the final
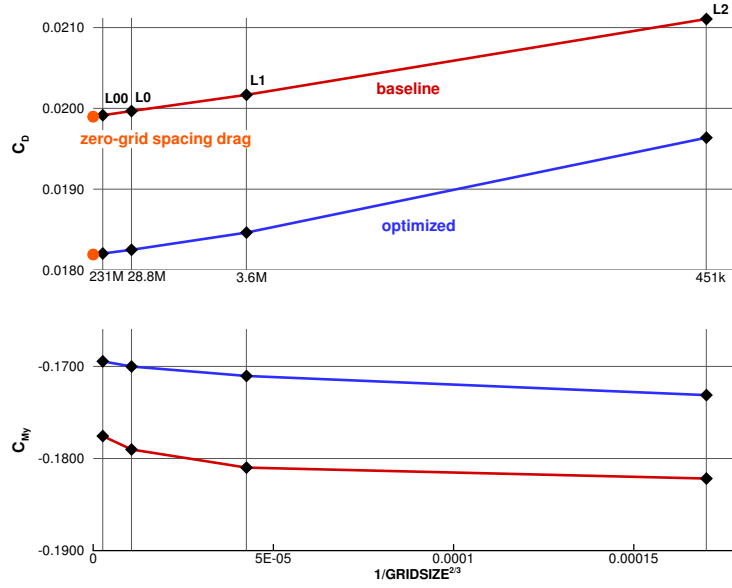
Figure 6: The mesh convergence study shows that the difference between the drag value computed with the 28.8 M grid and the zero-grid spacing drag is within one count.

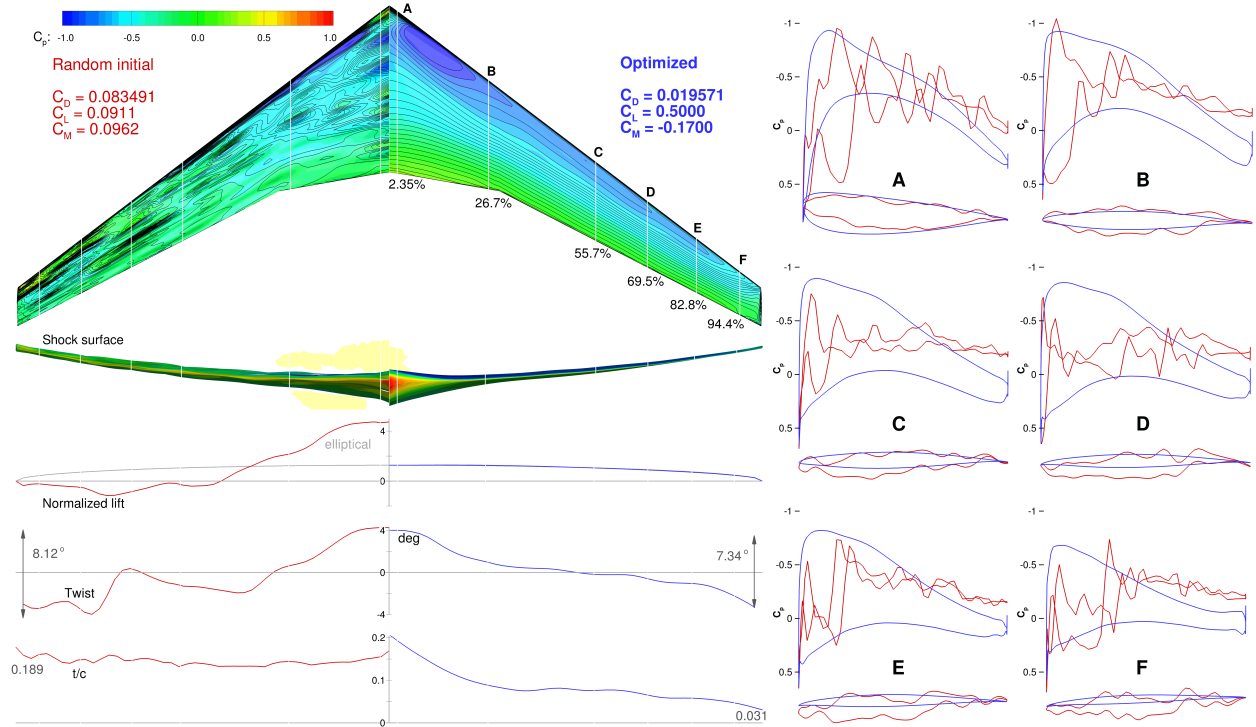result was the same as in Fig. 5.

Figure 7: The optimization manages to start from a random geometry and converge to an optimal wing that is shock free.

11

Lyu et al. (2015) present many other results for other problems related to this benchmark, including multipoint optimizations. Further multipoint optimization studies are performed by Kenway and Martins (2016b).

3.2 Full configuration with buffet constraints

For transonic wing design, aerodynamic shape optimization has demonstrated the potential produce high performance designs, but these designs are susceptible to buffet. To address this issue, we present a separation-based constraint formulation that can be used to constrain buffet onset in an aerodynamic shape optimization (Kenway and Martins, 2016a). A series of optimizations based on the AIAA ADODG wing-body-tail case are presented to show the need for buffet-onset constraints, and demonstrate the effectiveness of the proposed approach.

To quantify the buffet onset in aerodynamic shape optimization, we develop a new prediction method based on the amount of separated flow present on the wing in a steady RANS. We then use this method to enforce a constraint on the amount 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 ensures that the design stays within the boundary defined by the 30% margin to buffet onset airworthiness requirement. The constraint function is smooth and its gradient with respect to the wing shape variables is computed used a discrete adjoint implementation. We demonstrate the effectiveness of our new approach on the multipoint drag minimization of the wing-body-tail geometry (Kenway and Martins, 2015a).

To develop a more direct way of constraining buffet onset, we analyze 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 of this type of separation with increasing angle of attack is shown in Fig. 8. The first row in Fig. 8 shows the surface streamlines and pressure coefficient, as well as the shock sensor in orange (Lovely and Haimes, 1999).

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 (approximately the negative x-axis) direction, i.e., if

$$\cos\theta = \frac{\vec{V} \cdot \vec{V}_\infty}{|\vec{V}||\vec{V}_\infty|} < 0, \tag{4}$$

where θ is the angle between the local surface velocity and the freestream. Then we can define a *separation sensor* as,

$$\chi = \begin{cases} 1 & \text{if} \quad \cos\theta \leq 0 \\ 0 & \text{if} \quad \cos\theta > 0 \end{cases} \tag{5}$$

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 red areas on the surface for $\alpha = 3°$ and $\alpha = 3.29°$ in the bottom row of Fig. 8 shows the regions where $\chi = 1$, and they are an approximation to areas where the flow is separated.
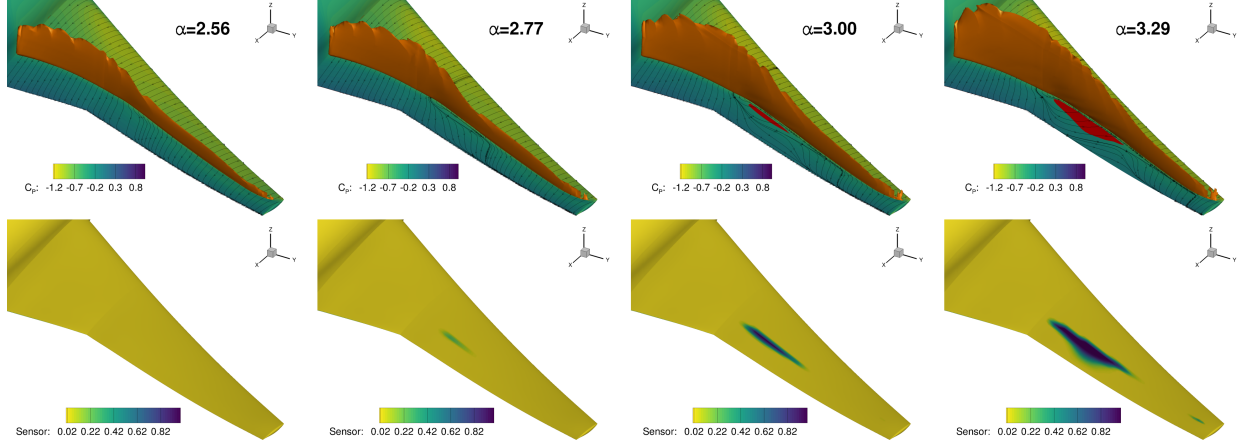
12

Figure 8: 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 reversed flow (red) and the shock (orange). Bottom row shows the separation area integrand value from Eqn. (7).

Our hypothesis is that we can correlate the value of the area where $\chi = 1$ to 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 an gradient-based optimization, we would like this function to be smooth, so we use a smooth Heaviside function to blend the discontinuity as follows,

$$\bar{\chi} = \frac{1}{1 + e^{2k(\cos\theta + \lambda)}} \tag{6}$$

In the 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.

Now, we can integrate the smooth separation sensor (6) over the surface and normalize it by the aircraft reference area to obtain the proposed *separation metric*:

$$S_{\text{sep}} = \frac{1}{S_{\text{ref}}} \int_S \bar{\chi} \, \mathrm{d}S. \tag{7}$$

This is equivalent to performing a weighted area integration of the sensor value shown in the bottom row of Fig. 8. We have shown that the separation metric (7) is correlated to buffet onset, by comparing to another method, as well as experimental data (Kenway and Martins, 2016a).

The baseline geometry defined in the ADODG Case 5 is taken directly from the 4^{th} Drag Prediction Workshop (Vassberg et al., 2014), and is shown in Fig. 9. 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 in satisfying buffet requirements. These cases—numbered 5.1 through 5.7—are summarized in Table 3. The objective of all optimizations (except for Case 5.7, which is

13

discussed separately) is to reduce the weighted drag coefficient at the N defined operating conditions. The optimization problem statement can be written as:

$$
\begin{array}{llll}
\text{minimize} & \sum_{i=1}^{N}\mathcal{W}_i C_{D_i} & \text{Quantity} \\[1em]
\text{with respect to} & \text{Wing cross sectional shape} & 240 \\
& \text{Wing twist} & 9 \\
& \text{Angle of attack } (\alpha_i) & N \\
& \text{Tail rotation angle } (\eta_i) & N \\
\text{subject to} & C_{L_i} - C_{L_i}^* = 0.0 & N \\
& C_{M_{y_i}} = 0.0 & N \\
& t_j \geq t_{j_{\text{CRM}}} & 750 \\
& S_{\text{sep}_i} \leq 0.04 & N
\end{array}
\tag{8}
$$

Only Cases 5.4, 5.5, and 5.7 use the separation constraint, where the constraint is applied only to the last two operating conditions. In practice, the adjoint for C_D is not evaluated for the buffet onset conditions, and conversely the separation metric adjoint is not evaluated for the normal operating conditions. This results in a total of three adjoint solutions required for each flight condition, which is desirable from a load balancing perspective. The constraints that the wing thickness may not be less than the baseline configuration at any location is particularly restrictive.

The ADODG specification for Case 5 requires the parameterization not to modify the planform, and any shape modification must be made only in the vertical direction. Additionally, twist rotation is permitted for the wing, as well as a solid-body rotation of the horizontal tail for trimming the aircraft. We use the FFD approach described in Sec. 2. The FFD volume and the associated geometric design variables are shown in Fig. 9.
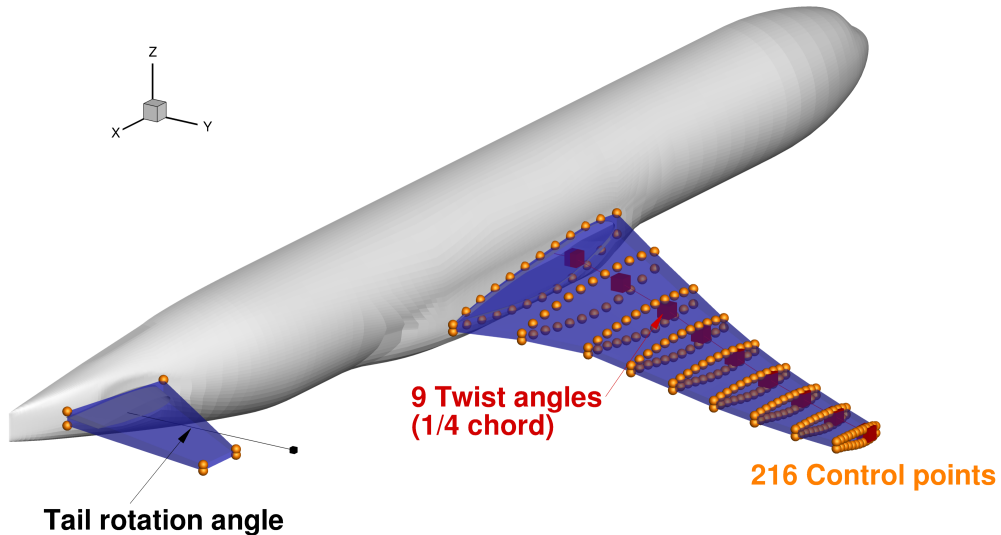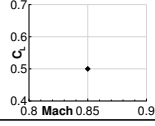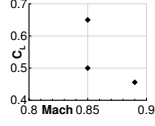
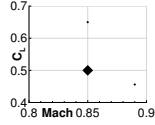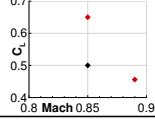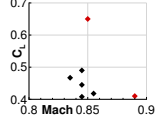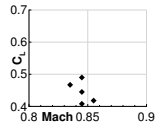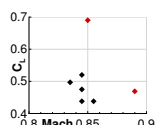Figure 9: Shape and twist modifications are permitted for the wing, and the horizontal tail can rotate.

Figures 10 and 11 show a summary of the key features of the two ADODG optimizations (Cases 5.1 and 5.2). The baseline configuration results are shown in red, while the optimized

14

Table 2: Operating conditions for each optimization case

Case	Point	Weights ($\mathcal{T}_i$)	Mach	C_L	Re	M–C_L plot
5.1	1	1	0.85	0.500	43.00×10^6	
5.2	1	$\frac{1}{3}$	0.85	0.500	43.00×10^6	
	2	$\frac{1}{3}$	0.85	0.650	43.00×10^6	
	3	$\frac{1}{3}$	0.89	0.456	45.00×10^6	
5.3	1	$\frac{2}{3}$	0.85	0.500	43.00×10^6	
	2	$\frac{1}{6}$	0.85	0.650	43.00×10^6	
	3	$\frac{1}{6}$	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	$\frac{1}{5}$	0.845	0.490	42.75×10^6	
	2	$\frac{1}{5}$	0.845	0.445	42.75×10^6	
	3	$\frac{1}{5}$	0.845	0.408	42.75×10^6	
	4	$\frac{1}{5}$	0.835	0.467	42.24×10^6	
	5	$\frac{1}{5}$	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	$\frac{1}{5}$	0.845	0.490	42.75×10^6	
	2	$\frac{1}{5}$	0.845	0.445	42.75×10^6	
	3	$\frac{1}{5}$	0.845	0.408	42.75×10^6	
	4	$\frac{1}{5}$	0.835	0.467	42.24×10^6	
	5	$\frac{1}{5}$	0.855	0.418	43.25×10^6	
5.7	1	$\frac{1}{5}$	0.845	0.520	42.75×10^6	
	2	$\frac{1}{5}$	0.845	0.475	42.75×10^6	
	3	$\frac{1}{5}$	0.845	0.438	42.75×10^6	
	4	$\frac{1}{5}$	0.835	0.497	42.24×10^6	
	5	$\frac{1}{5}$	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	

results are shown in blue. 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 increased by 2.8 counts. The drag divergence curves indicate a significant drag penalty across the lower Mach numbers but it does have a much higher drag divergence Mach number than the baseline design.
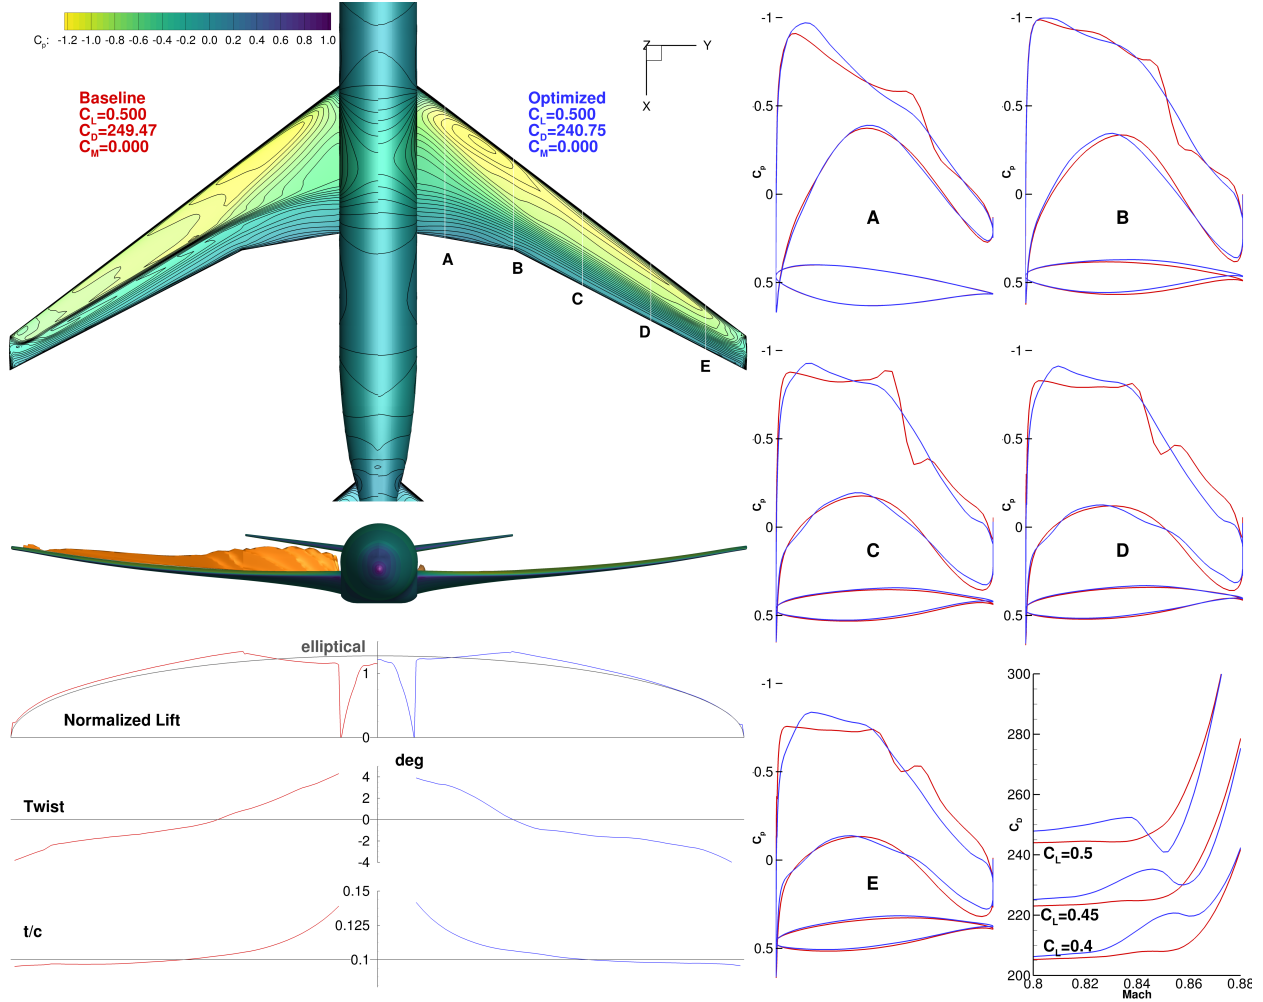
Figure 10: High performance is achieved for the single point design at the nominal operating condition.

While drag coefficient divergence curves yield useful insights into optimized designs, it is particularly instructive to examine the 2D analogue in the form of ML/D contours. In the context of transonic transport wing design, ML/D is a better measure of performance, since it includes the benefit that flying faster has on the overall aircraft efficiency.

Contour plots for all seven optimizations including the baseline design are shown in Fig. 12. The contours in each figure extend up to the predicted buffet onset curve shown in red. The orange curve shows the prediction buffet onset using the $\Delta\alpha = 0.1$ method. Overall, the separation metric method continues to compare well with 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 from the red buffet onset curve. Only operating conditions below the buffet margin curve can be considered for normal operation. The absolute maximum ML/D value for

16
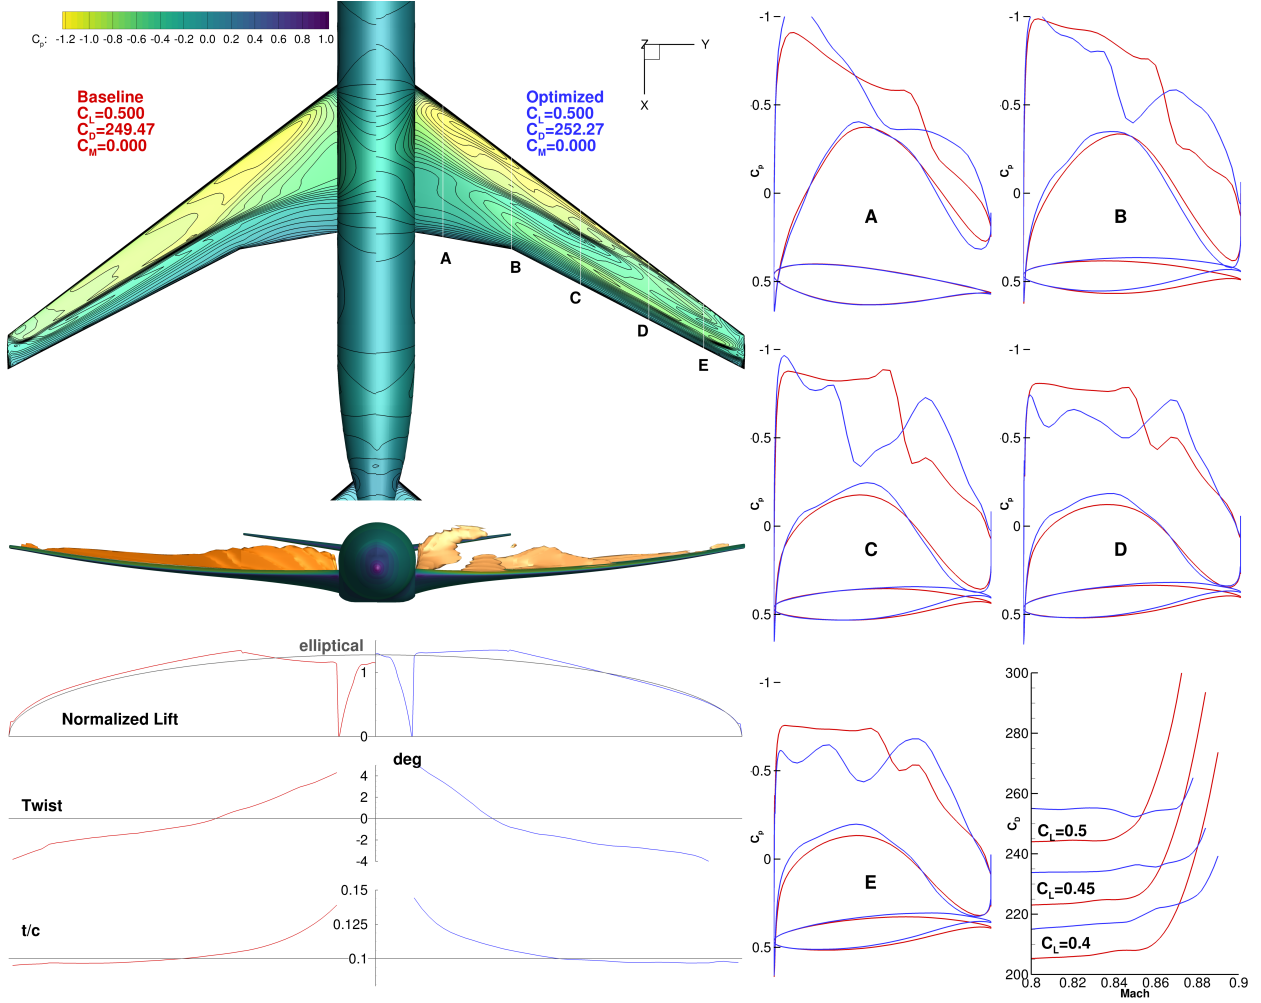
Figure 11: Performance is sacrificed across a large range of Mach numbers to obtain a small amount of improvement at the highest Mach numbers.

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% $(ML/D)_{\max}$ of 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 also decreases in the flight time. The area enclosed by both of these contours is used to quantify the design's robustness in these figures. The areas are scaled by a factor of 100^2 such that the area of the rectangle measuring 0.01 in M and 0.01 in C_L has an area of 1.

The design operating conditions listed in Table 2 are shown as diamonds. The operating conditions considered for the objective function are shown in black, while the buffet onset constraint conditions are shown in red. The first point ($M = 0.85$, $C_L = 0.65$) is at the nominal cruise Mach number and the C_L value corresponding to a 1.3 g 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})

17

condition. The lift coefficient for this condition is adjusted to give the same dimensional lift as the nominal cruise condition.

There are two additional regions that are highlighted in black and orange. We refer to these 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 like 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 4 000 ft decrease in altitude. To put it in another way, the integration region contains all operating conditions within 4 000 ft of the buffet-constrained ceiling and at all normal operating Mach numbers.

The contour plots give a much more complete understanding of the optimized designs. Unsurprisingly, the single point optimization—Case 5.1—was able to produce the highest ML/D value, which is almost exactly matched to the design operating condition. However, without a way to constrain the buffet onset boundary, the $(ML/D)_{\text{max}}$ 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% $(ML/D)_{\text{max}}$ contour (blue) is small, indicating a highly localized point-design. Despite the highest ML/D value, the average ML/D in its own integration region (orange) is 4.6% worse than the baseline design.

For Case 5.2, the addition of operating conditions at the edge of the buffet onset envelope has the effect of substantially improving the buffet boundary over the entire range of Mach numbers. This case resulted in the most robust buffet onset behavior of all cases. However, the $(ML/D)_{\text{max}}$ has barely improved over the baseline design (17.18 vs. 17.13). Worse, 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 ML/D average in the integration region is slightly worse (-0.5%) than 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 value obtained by integration of 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 weight of 2/3, while the remaining two points have weights of 1/6 each. The adjusted weightings yield a much more useful design. This is the first case where a significant portion of 99% $(ML/D)_{\text{max}}$ contour falls within the integration region. Additionally, 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 is able to operate at higher altitudes. The other issue with this case is that the specific weightings were arbitrarily picked. These particular weights yielded acceptable results, but there is no guarantee these weight values would 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
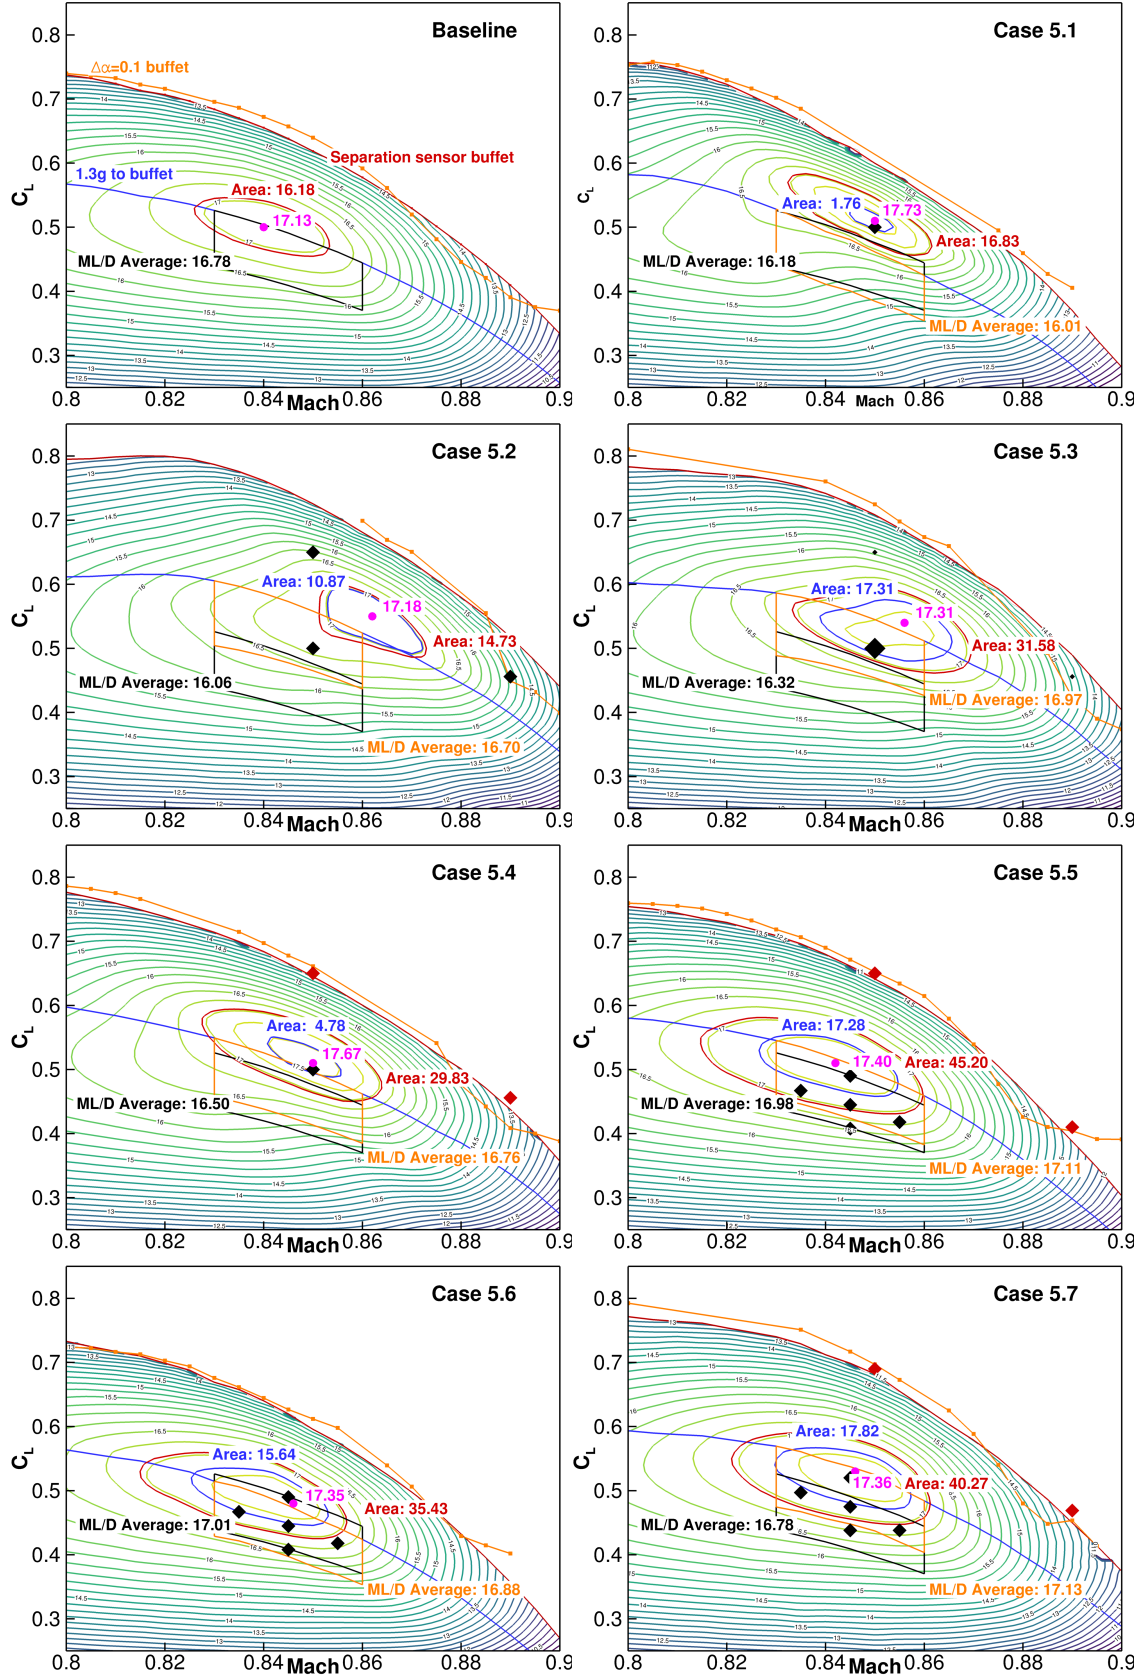
Figure 12: Contours of ML/D for the baseline and each optimized configuration.

sensor and constrain its value. Note that there is a slight discrepancy between the operating conditions (red diamonds) and the buffet onset boundary itself. The reason is that the buffet onset conditions are analyzed using the JST scalar dissipation scheme, which results in a solution with more dissipation. The scalar scheme provides the increased robustness necessary for optimization, which is unnecessary for the contour plot evaluations. The more dissipative scalar scheme slightly under-predicts the area of separated flow, and thus the buffet boundary is lower when analyzed with the matrix scheme for the contour plot. Overall, the performance of this design is similar to that of the single-point optimization (Case 5.1). Most of the high performance region lies outside the integration region. However, the performance reduction is not as pronounced as with Case 5.1, with a 1.6% performance reduction in the original integration region and almost zero for the on-design integration region. However, there is a small improvement in the buffet onset boundary.

After analyzing the results from Cases 5.1–5.4, we noticed that the nominal design point was always located towards the upper side or even completely outside the integration region. It is perhaps no surprise that all optimization discussed thus far failed to improve the performance in the baseline integration region. To address this issue, we formulated a multipoint optimization, Case 5.5. Previous optimizations performed by the authors on the CRM wing along configuration have shown that five operating conditions arranged as a cross in M–C_L space results is highly robust designs (Kenway and Martins, 2016b). Since our goal is to improve the performance in the original integration region, we distribute five operating conditions as follows. The nominal Mach number is reduced to 0.845 for the first three operating conditions. The first point lies on the 1.3 g buffet onset boundary, while the second two points are at C_L values corresponding to 2 000 and 4 000 ft below. The two remaining points are 2 000 ft lower than the buffet onset boundary at ± 0.01 Mach number. The buffet onset conditions are taken at $M = 0.85$, $C_L = 0.65$ and $M = 0.89$, $C_L = 0.41$. The later 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 has now increased by 1.2%, the performance of the updated integration region increase by 2.0%. The design is very robust, as shown by area inside the 99% $(ML/D)_{\max}$ contour. Further, the maximum performance point is located inside the operating envelope.

Next we developed Case 5.6 to investigate the effect of removing the buffet onset conditions present in Case 5.5. We wish to answer the 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 dropped slightly over the integration envelope, pushing the integration region into a lower performance region. The average ML/D for the integration region is 16.88, only 0.6% higher than the baseline configuration. This is much smaller than the 2.0% improvement obtained in Case 5.5.

For the last case, 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 C_L for the buffet onset locations. In the formulation of Case 5.7, we want the optimization to adjust the single nominal operating condition directly. 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. The

high Mach buffet case must have the same physical lift as the nominal operating condition at $M = 0.89$. Finally, the remaining operating conditions move vertically in sync with the changing design C_L. The modified optimization formulation is as follows:

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

Note that operating conditions (diamonds) shown in Fig. 12 are the optimized values. The optimization increased the nominal design C_L from the initial value of 0.490 (the same value as used in Cases 5.5 and 5.6) and increased it to 0.520, an increase of 0.03 in C_L. This increase is only possible due to 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% $(ML/D)_{\text{max}}$ contour of all cases.

4 Aerostructural optimization framework

Coupling aerodynamic and structural numerical models to compute the static aeroelastic shape of lifting surfaces is essential when designing lifting surfaces that are flexible. Even small changes in shape can have a large effect on the aerodynamic performance, and multiple flow conditions result in multiple shapes. This is particularly important for swept wings, where bend-twist coupling can result in large changes in the twist distribution. In this section we extend the aerodynamic shape optimization framework by coupling it to a structural solver, and demonstrate the developed approach in two wing design problems.

4.1 Structural solver and adjoint

The structural solver used in this work is the Toolkit for the Analysis of Composite Structures (TACS) (Kennedy and Martins, 2014). TACS includes an adjoint solver that is able to handle the structural design variables, which in our case are the thicknesses of the structural members. Parallelism is achieved within TACS by using an element-based partitioning of the finite-element mesh. This partitioning is used to parallelize the factorization of the stiffness matrix, the computation and assembly of the stiffness matrix and structural residuals, and the computation of the functions of interest and their derivatives. Typically, the

factorization of the stiffness matrix is the most costly operation. To parallelize the matrix factorization and back solutions, TACS uses a Schur-complement-based parallel direct solver. In this technique, each processor independently computes the local contribution to a reduced linear system that is formed from the all unknowns on the domain interface. This reduced problem is the global Schur complement. In TACS, the local contributions to the global Schur complement are computed using a block-based parallel factorization. After the global Schur complement is computed, it is factored in parallel using a sparse block-cyclic algorithm that achieves excellent parallel performance. This direct method enables us to solve poorly conditioned thin-shell structural problems with condition numbers $\mathcal{O}(10^9)$ in an efficient manner. For general nonlinear analysis the structural discipline residuals are

$$\mathcal{S}(u) = 0 \tag{10}$$

where u is the vector of structural displacements. For linear analysis, this equation can be written as $\mathcal{S}(u) = Ku - F$, where K is the linear stiffness matrix and F is the load vector.

4.2 Load and displacement transfer

Our load and displacement transfer scheme follows the work of Brown (1997). In this approach, rigid links are used to extrapolate the displacements from the structural surface to the CFD surface, as shown in Fig. 13. These rigid links are constructed between the aerodynamic surface mesh points and the points on the structural model lying closest to this set of points. 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 is performed on the aerodynamic mesh and is transmitted back through the rigid links to the structure.
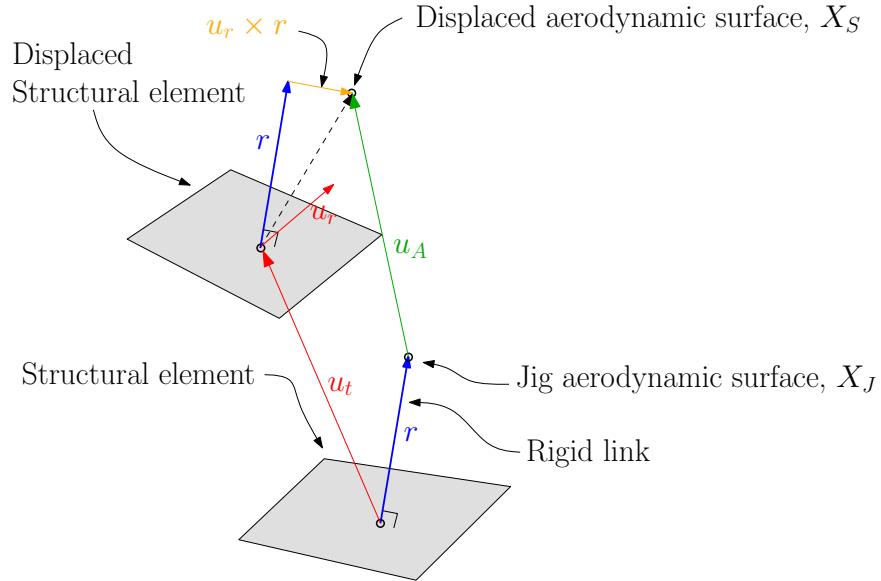

Figure 13: Load-displacement transfer operation

The design jig shape, X_J, is uniquely determined by the optimization design variables x. A single coordinate on the perturbed surface, X_S, is given by

$$X_S = X_J + u_A = X_J + u_t + u_r \times r \tag{11}$$

where r is the vector that connects a point on the aerodynamic surface to the closest location on the structural mesh, u_t is the translational component of the displacement, and u_r is the small-angle approximation of the rotations in the global reference frame, as shown in Fig. 13. We can compactly represent the displacement transfer from u to u_A as

$$u_A = Tu = \left(\frac{\partial u_A}{\partial u} \right) u \tag{12}$$

where T is the generalized transfer matrix. To transfer loads from the aerodynamic discipline to the structural discipline, we employ the transpose operation,

$$F = T^T F_A = \left(\frac{\partial F}{\partial F_A} \right) F_A \tag{13}$$

where F_A are the forces on the CFD nodes. More details on these transfers can be found in Brown (1997) and Martins et al. (2005).

4.3 Coupled solution algorithms

We formulate the aerostructural analysis problem using a two-field formulation. The aerodynamic analysis is the first field and the structural analysis is the second. Using this formulation, we can write the governing equations of both disciplines as a function of the fluid states, w, structural states, u, and design variables x. The latter is a vector consisting of global variables and local variables. Global variables affect the two disciplines directly, while local variables affect only a single discipline in a direct manner. Geometric variables that change the aircraft wetted surface—or outer mold line (OML)—are global variables. For example, an airfoil thickness variable will not only change the aerodynamic shape but also the height of the ribs and spars in the internal structure. Local variables include the angle of attack, which directly affects only the aerodynamics, and the thicknesses of the structural members, which affect only the structures. Combining the residual equations from the aerodynamic and structural disciplines, we write the combined residual of the multidisciplinary system as

$$\mathcal{R} = \begin{bmatrix} \mathcal{A}(w, u; x) \\ \mathcal{S}(u, w; x) \end{bmatrix} = 0. \tag{14}$$

The aerostructural analysis consists in finding a solution, (w, u), that satisfies these coupled residual equations.

4.3.1 Nonlinear block Gauss–Seidel method (NLBGS)

The traditional process for solving the coupled aerostructural equations (14) is to use a nonlinear block Gauss–Seidel (NLBGS) method (Martins et al., 2005; Maute et al., 2001). In this approach, the aerodynamic analysis is first partially converged and the aerodynamic forces are evaluated. The forces are then transferred to the structural analysis and the corresponding displacements are computed. Finally, the displacements are transferred back to the aerodynamic analysis, the mesh is deformed, a new CFD solution is found, and this iterative loop continues until the coupled convergence criterion is met. The NLBGS procedure is listed in Algorithm 1.

Algorithm 1 Nonlinear block Gauss–Seidel method

1: **Given**: $w^{(0)}, u^{(0)}, X_J, k_{\max}$

2: **Initialize**: $\theta \in (0, 1]$

3: **for** $k \leftarrow 1, k_{\max}$ **do**

4:    $X_S^{(k)} \leftarrow Tu^{(k-1)} + X_J$ ▷ Transfer displacements from CSM to CFD

5:    $X_V^{(k)} \leftarrow \mathcal{W}\left(X_S^{(k)}\right)$ ▷ Deform volume mesh to match surface

6:    $\mathcal{A}^{(k)} \leftarrow \mathcal{A}\left(w^{(k-1)}, X_V^{(k)}\right)$ ▷ Evaluate initial CFD residual

7:    **while** $\|\mathcal{A}\left(w_i; x\right)\| > \epsilon_{\mathcal{A}} \left\|\mathcal{A}^{(k)}\right\|$ **do**   ▷ CFD partial convergence criterion

8:        $w_{i+1}^{(k-1)} \leftarrow w_i^{(k-1)} + \Delta w_i$ ▷ Iterate CFD

9:        $i \leftarrow i + 1$

10:   **end while**

11:   $w^{(k)} \leftarrow w_i^{(k-1)}$ ▷ New CFD solution

12:   $F_A^{(k)} \leftarrow F_A\left(w^{(k)}, X_S^{(k)}\right)$ ▷ Evaluate aerodynamics forces

13:   $F^{(k)} \leftarrow T^T F_A^{(k)}$ ▷ Transfer forces from CFD to CSM

14:   $\mathcal{S}^{(k)} \leftarrow \mathcal{S}\left(u^{(k-1)}, F^{(k)}\right)$ ▷ Evaluate initial CSM residual

15:   **while** $\left\|\mathcal{S}\left(u^{(k)}; x\right)\right\| > \epsilon_{\mathcal{S}} \left\|\mathcal{S}^{(k)}\right\|$ **do**   ▷ CSM partial convergence criterion

16:        $u_{i+1}^{(k-1)} \leftarrow u_i^{(k-1)} + \Delta u_i$ ▷ Iterate CSM

17:        $i \leftarrow i + 1$

18:   **end while**

19:   $u^{(k)} \leftarrow u_f^{(k-1)}$ ▷ Temporary new CSM solution

20:   $\Delta u^{(k)} \leftarrow u^{(k)} - u^{(k-1)}$ ▷ Compute displacement increment

21:   **if** $k > 1$ **then**

22:        $\theta \leftarrow \theta\left(1 - \dfrac{\left(\Delta u^{(k)} - \Delta u^{(k-1)}\right) \cdot \Delta u^{(k)}}{\left\|\left(\Delta u^{(k)} - \Delta u^{(k-1)}\right)\right\|^2}\right)$ ▷ Adapt under-relaxation with Aitken acceleration

23:   **end if**

24:   $u^{(k+1)} \leftarrow u^{(k)} + \theta \Delta u^{(k)}$ ▷ Compute new structural states

25:   **if** $\left\|\mathcal{A}^{(k)}\right\| < \epsilon_{AS} \left\|\mathcal{A}^{(1)}\right\|$ **and** $\|\mathcal{S}^{(k)}\| < \epsilon_{AS}\|\mathcal{S}^{(1)}\|$ **then**   ▷ Aerostructural convergence criterion

26:        **break**

27:   **end if**

28: **end for**

Three tolerances are defined for the coupled analysis. The relative tolerances, $\epsilon_{\mathcal{A}}$ and $\epsilon_{\mathcal{S}}$, are the tolerances required by the aerodynamic and structural disciplines, respectively, for each NLBGS iteration. The aerodynamic solver tolerance is typically $\mathcal{O}\left(10^{-1}\right)$, while $\epsilon_{\mathcal{S}}$ is typically $\mathcal{O}\left(10^{-3}\right)$ or smaller. The third tolerance, ϵ_{AS}, is the aerostructural solution tolerance and represents the feasibility of the interdisciplinary coupling. Typical values for ϵ_{AS} range from 10^{-3} for an engineering solution accurate to three decimal places to 10^{-6},

which is the typical value used when performing design optimization. Aitken acceleration (Irons and Tuck, 1969) (Line 22 of Algorithm 1) is employed to dynamically choose the under-relaxation factor to accelerate convergence. One advantage of this method is that each disciplinary solver can be used without modification. For tightly coupled aerostructural problems with large displacements, however, this method may converge slowly or not at all.

4.3.2 Coupled Newton–Krylov method (CNK)

The second approach that we use for the aerostructural solution is a fully coupled Newton–Krylov (CNK) method. Algorithm 2 lists the pseudocode for evaluating the coupled nonlinear residual, $\mathcal{R}$.

Algorithm 2 Coupled nonlinear residual computation

1: **function** $\mathcal{R}(w,u)$
2:    $X_S \leftarrow Tu + X_J$ ▷ Transfer displacements
3:    $X_V \leftarrow \mathcal{W}(X_S)$ ▷ Deform volume mesh to match surface
4:    $\mathcal{A} \leftarrow \mathcal{A}(w, X_V)$ ▷ Evaluate CFD residuals
5:    $F_A \leftarrow F_A(w, X_S)$ ▷ Evaluate aerodynamics forces
6:    $F \leftarrow T^T F_A$ ▷ Transfer forces
7:    $\mathcal{S} \leftarrow \mathcal{S}(u, F)$ ▷ Evaluate CSM residuals
8:    $\mathcal{R} \leftarrow (\mathcal{A}, \mathcal{S})$ ▷ Combine residuals
9: **return** $\mathcal{R}$
10: **end function**

The procedure is similar to one iteration of the NLBGS method, except that instead of computing an approximate solution update to the state variables, we evaluate only residuals. We use an inexact Newton–Krylov approach to solve the coupled equations (14) by computing the approximate Newton update,

$$\begin{bmatrix} \dfrac{\partial \mathcal{A}}{\partial w} & \dfrac{\partial \mathcal{A}}{\partial u} \\ \dfrac{\partial \mathcal{S}}{\partial w} & \dfrac{\partial \mathcal{S}}{\partial u} \end{bmatrix} \begin{bmatrix} \Delta w \\ \Delta u \end{bmatrix} = - \begin{bmatrix} \mathcal{A}(w) \\ \mathcal{S}(u) \end{bmatrix}. \tag{15}$$

After this update is computed, the state variables w and u are updated with Δw and Δu, respectively.

To reduce the memory requirement, we use a matrix-free method. The matrix-vector products required by the Krylov method are approximated with finite differences. In theory, with this approach, we need only residual evaluations to solve the coupled problem. However, Krylov methods require effective preconditioning for acceptable performance, especially for very large systems with millions of degrees of freedom. We have implemented a block-Jacobi preconditioner that reuses the preconditioner of the original discipline solvers. This has two advantages: 1) the preconditioners for each discipline can be applied in parallel, because this approach eliminates contributions from the off-diagonal terms, which are never explicitly stored, and 2) since both of our discipline solvers use implicit methods, this preconditioner simply reuses the same linear solution method from the original solvers. Since the coupled

preconditioning method varies from iteration to iteration, we must use a flexible variant of a Krylov method to solve the coupled Newton update (15). We choose to use FGMRES (Saad and Schultz, 1986), which has been shown to work well for a wide range of large asymmetric systems.

4.4 Coupled adjoint derivatives

The critical component of the aerostructural framework is an efficient method for the computing coupled derivatives. Using the approach outlined by Martins et al. (2005), we write the adjoint equations for the coupled aerostructural system. This is a specific case of the unifying chain rule presented in Part 1 of this lecture. The residuals, state variables, and adjoint variables for the aerodynamic discipline are $\mathcal{A}$, w, and ψ, respectively, and the corresponding variables for the structural discipline are $\mathcal{S}$, u, and ϕ. The total derivative of the function of interest, f, is

$$\mathrm{d}fx = \frac{\partial f}{\partial x} + \begin{bmatrix} \frac{\partial f}{\partial w} & \frac{\partial f}{\partial u} \end{bmatrix} \begin{bmatrix} \mathrm{d}wx \\ \mathrm{d}ux \end{bmatrix}. \tag{16}$$

We write the total derivative of the residuals as

$$\begin{bmatrix} \mathrm{d}\mathcal{A}x \\ \mathrm{d}\mathcal{S}x \end{bmatrix} = \begin{bmatrix} \frac{\partial \mathcal{A}}{\partial x} \\ \frac{\partial \mathcal{S}}{\partial x} \end{bmatrix} + \begin{bmatrix} \frac{\partial \mathcal{A}}{\partial w} & \frac{\partial \mathcal{A}}{\partial u} \\ \frac{\partial \mathcal{S}}{\partial w} & \frac{\partial \mathcal{S}}{\partial u} \end{bmatrix} \begin{bmatrix} \mathrm{d}wx \\ \mathrm{d}ux \end{bmatrix} = 0. \tag{17}$$

Substituting the solution of Eq. (17) into Eq. (16) to eliminate the total derivatives we obtain

$$\mathrm{d}fx = \frac{\partial f}{\partial x} - \underbrace{\begin{bmatrix} \frac{\partial f}{\partial w} & \frac{\partial f}{\partial u} \end{bmatrix} \begin{bmatrix} \frac{\partial \mathcal{A}}{\partial w} & \frac{\partial \mathcal{A}}{\partial u} \\ \frac{\partial \mathcal{S}}{\partial w} & \frac{\partial \mathcal{S}}{\partial u} \end{bmatrix}^{-1}}_{\Psi^T} \begin{bmatrix} \frac{\partial \mathcal{A}}{\partial x} \\ \frac{\partial \mathcal{S}}{\partial x} \end{bmatrix}. \tag{18}$$

This yields the coupled adjoint equations,

$$\begin{bmatrix} \frac{\partial \mathcal{A}}{\partial w} & \frac{\partial \mathcal{A}}{\partial u} \\ \frac{\partial \mathcal{S}}{\partial w} & \frac{\partial \mathcal{S}}{\partial u} \end{bmatrix}^T \begin{bmatrix} \psi \\ \phi \end{bmatrix} = \begin{bmatrix} \frac{\partial f}{\partial w} & \frac{\partial f}{\partial u} \end{bmatrix}^T. \tag{19}$$

After the solution for the coupled adjoint equations (19) is obtained, the following equation can be used to compute the total derivative:

$$\mathrm{d}fx = \frac{\partial f}{\partial x} - \psi^T \left(\frac{\partial \mathcal{A}}{\partial x} \right) - \phi^T \left(\frac{\partial \mathcal{S}}{\partial x} \right). \tag{20}$$

Implementing the coupled adjoint, including all the required partial-derivative terms, is a challenging endeavor. Furthermore, ensuring that the partial-derivative computations and solution methods are efficient and exhibit good parallel scalability is even more difficult.

26

The aerodynamic diagonal block, $\partial \mathcal{A}/\partial w$, in the coupled adjoint equations (19), represents the Jacobian that we already explained in Sec. 2.

The first off-diagonal block, $\partial \mathcal{A}/\partial u$, contains the derivatives of the aerodynamic residual with respect to the structural displacements. This is a challenging matrix to compute in a two-state aerostructural formulation. What makes this term particularly challenging is that a single structural degree of freedom (DOF) can affect many aerodynamic cells. Using the mesh deformation described in Sec. 2.3, any structural DOF that perturbs a CFD node results in a dense column due to the solution of the linear system of equations. With a large number of structural surface DOFs and a large number of CFD nodes, storing this matrix would require an excessive amount of memory. Instead, we compute this term in a matrix-free fashion using the chain rule:

$$\left(\frac{\partial \mathcal{A}}{\partial u}\right)^T \psi = \left(\frac{\partial X_S}{\partial u}\right)^T \left(\frac{\partial X_V}{\partial X_S}\right)^T \left(\frac{\partial \mathcal{A}}{\partial X_V}\right)^T \psi \tag{21}$$

where $\partial \mathcal{A}/\partial X_V$ contains the derivatives of the aerodynamic residual with respect to variations of all the volume mesh coordinates. The sparse structure of this matrix is due to the stencil used for the finite-volume computation. We compute this matrix using reverse-mode AD, in a similar fashion to $\partial \mathcal{A}/\partial w$. This matrix is computed in parallel and stored in distributed memory, which allows us to compute the transpose matrix-vector products in a relatively inexpensive manner.

The matrix of derivatives of the mesh deformation, $\partial X_V/\partial X_S$, requires a careful implementation to be computationally efficient. A summary of the full mesh-deformation procedure is shown in Fig. 14.

$$X_S \xrightarrow{\;Y_S = RX_S\;} Y_S \xrightarrow{\;K_{uu}Y_V - K_{us}Y_S = 0\;} Y_V \xrightarrow[\text{interpolation}]{\text{Bézier}} X_{SW} \longrightarrow X_V$$
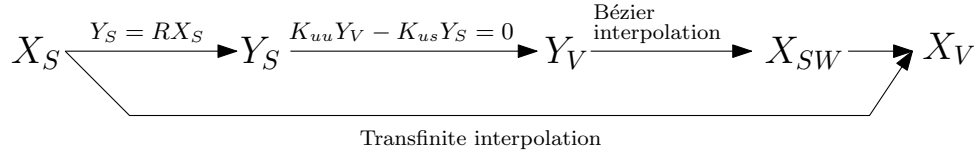
$$\text{Transfinite interpolation}$$

Figure 14: Solid mesh warping procedure

For the transpose matrix-vector products required for the adjoint computation, the normal deformation procedure of Fig. 14 is done in reverse. Reverse-mode AD is used to compute the variation of all the surface nodes $(\overline{X_S})$ and the variation of the volume supernodes $(\overline{Y_V})$. The variation of the volume supernodes is related to the surface supernodes through the solution of the mesh adjoint equation given by

$$K_{uu}^T \psi_M = -\overline{Y_V}, \tag{22}$$

where ψ_M is the mesh adjoint. Finally, the variation $\overline{Y_S}$ is propagated to the full mesh surface using the transpose of the restriction operator. This procedure is illustrated in Fig. 15.

This computation is performed just once for each transpose matrix-vector product. To ensure overall scalability of the coupled adjoint solution, all of the computations are executed in parallel with acceptable efficiency. In this case, since the linear system of equations for the mesh deformation is symmetric and fixed, a parallel LU decomposition is performed
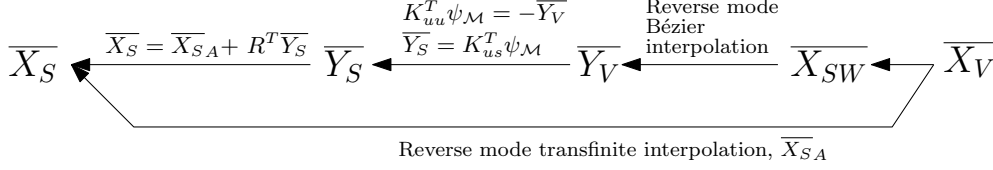
27

Figure 15: Adjoint procedure for solid mesh warping

once using the distributed version of SuperLU. Subsequent mesh perturbations and adjoint solutions can be computed quickly with a back-substitution operation. Finally, since X_S is given by $X_S = X_J + Tu$, the partial-derivative matrix $(\partial X_S/\partial u)^T$ is simply T^T, which is the load transfer operation.

The final aerodynamic residual partial derivative, $\partial \mathcal{A}/\partial x$, is computed using another chain rule:

$$\psi^T \left(\frac{\partial \mathcal{A}}{\partial x} \right) = \psi^T \left(\frac{\partial \mathcal{A}}{\partial X_V} \right) \left(\frac{\partial X_V}{\partial X_S} \right) \left(\frac{\partial X_S}{\partial x} \right). \tag{23}$$

The two right-most terms are identical to the terms in the computation of $\partial \mathcal{A}/\partial u$ (21),

The other off-diagonal block in the coupled adjoint equations (19), $\partial \mathcal{S}/\partial w$, represents the derivatives of the structural residuals with respect to the aerodynamic states. The only contribution to this matrix is due to the aerodynamic forces, F_A, and is given by,

$$
\begin{aligned}
\left(\frac{\partial \mathcal{S}}{\partial w} \right)^T \phi &= - \left(\frac{\partial F}{\partial w} \right)^T \phi \\
&= - \left(\frac{\partial F_A}{\partial w} \right)^T \left(\frac{\partial F}{\partial F_A} \right)^T \phi \\
&= - \left(\frac{\partial F_A}{\partial w} \right)^T T\phi.
\end{aligned}
\tag{24}
$$

We compute and store $\partial F_A/\partial w$ and $\partial F_A/\partial X_S$, which we refer to as the *coupling matrices.* These terms are computed using reverse-mode AD applied to the CFD force evaluation routine. Unlike $\partial \mathcal{A}/\partial w$ or $\partial \mathcal{A}/\partial X_V$, only cells on the wetted surface have nonzero components. Thus, the cost of storing these terms is very low.

The remaining diagonal block, $\partial \mathcal{S}/\partial u$, represents the derivative of the structural residuals with respect to the structural states. For linear aerostructural analysis, this term is *not* simply the linear stiffness matrix, K. Given the structural residuals (10), we find that the applied forces, F, are actually explicit functions of u. As the flexible structure deforms, the surface normals on the CFD mesh change orientation, giving rise to a nonlinear following force. To obtain derivatives, this effect must be included. To obtain the matrix-vector products required by the Krylov-type solution strategy, we compute this term in the following

manner:

$$\left(\frac{\partial \mathcal{S}}{\partial u}\right)^T \phi = \left(K - \left(\frac{\partial F}{\partial u}\right)^T\right)\phi \tag{25}$$

$$= \left(K - \left(\frac{\partial F}{\partial F_A}\right)\left(\frac{\partial F_A}{\partial X_S}\right)^T\left(\frac{\partial X_S}{\partial u}\right)\right)\phi$$

$$= \left(K - T^T\left(\frac{\partial F_A}{\partial X_S}\right)^T T\right)\phi.$$

This product is formed by transferring the ϕ vector to the aerodynamics using a displacement transfer, finding the $\partial F_A/\partial X_S$ transpose matrix-vector product, and then transferring the result back to the structural solver using a load transfer. Since this operation requires two synchronous data transfers, these products are evaluated only as required, and the exact $\partial \mathcal{S}/\partial u$ is never formed.

The partial derivatives of the structural residuals with respect to the design variables, $\partial \mathcal{S}/\partial x$, required in total-derivative equation (18), are evaluated analytically.

The product of the partial derivatives and the adjoint vector $\psi^T \partial S/\partial x$ is evaluated analytically on an element-by-element basis, where geometric and material variables are handled using different approaches. For the geometric variables, the computation proceeds in two stages. First, the product $\psi^T \partial \mathcal{S}/\partial X_J$ is computed, where X_J are all the node locations. Next, the product of $\psi^T \partial \mathcal{S}/\partial X_J$ and the derivative of the node locations with respect to the design variables, $\partial X_J/\partial x$, is computed. This computation is arranged so that it is efficient when the number of geometric design variables is large. For the derivative of the material design variables, the sparse dependency of the design variables is exploited, and only elements with a nonzero contribution to $\psi^T \partial \mathcal{S}/\partial x$ are computed.

The right-hand side of the coupled adjoint system (19) is the derivative of the function of interest with respect to the system states. We consider only aerodynamic or structural functions, that is, functions that could be used in a single-discipline analysis. The derivatives of typical aerodynamic functions such as C_L, C_D, and C_M are computed with reverse-mode AD, in a similar fashion to the coupling matrices. We analytically evaluate the derivatives of typical structural functions, such as individual element stresses, σ_i, or the Kreisselmeier–Steinhauser (KS) aggregation functions (Poon and Martins, 2007; Wrenn, 1989). For the structural functions, $\partial I/\partial w$ is zero. However, for aerodynamic functions involving surface pressure or traction integration, $\partial I/\partial u$ is nonzero and is given by

$$\frac{\partial f}{\partial u} = \left(\frac{\partial f}{\partial X_S}\right)\left(\frac{\partial X_S}{\partial u}\right) \tag{26}$$

$$= \left(\frac{\partial f}{\partial X_S}\right)T.$$

The resulting procedure is similar to that used in the computation of $\partial \mathcal{S}/\partial u$, expressed in Eq. (25).

The evaluation of the partial-derivative terms in the total-derivative equation (20) is challenging because the load and displacement transfer operations have a dependency on the

geometric design variables. Figure 13 shows how the displaced surface coordinates, X_S, are extrapolated from the structural domain. The complication arises from the fact that the length of the extrapolation vector, r, changes with the design variables.

For aerodynamic functions, we compute $\partial I / \partial X_S$ using reverse-mode AD, similarly to what was done for the coupling matrices in Eq. (24). Then, for each design variable, we evaluate the derivative of the perturbed shape, X_S, with respect to the design variables, x, by computing

$$\frac{\partial f}{\partial x} = \frac{\partial f}{\partial X_S} \left(\frac{\partial X_S}{\partial x} \right) \tag{27}$$
$$= \frac{\partial f}{\partial X_S} \left(\frac{\partial X_J}{\partial x} + \frac{\partial (u_r \times r)}{\partial x} \right).$$

Since the transfer-operation matrix T is never explicitly formed, we evaluate a single entry of the resulting vector sequentially in a matrix-free fashion. A similar correction must be made for the derivative of the structural states with respect to the geometric design variables, since the load transfer depends on the design variables.

Finally, the evaluation of the $\partial X_S / \partial x$ term is used to complete the computation of $\partial \mathcal{A} / \partial x$. These operations are relatively costly, since in our implementation they require a synchronous transfer of data from both disciplines for each design variable.

4.5 Coupled adjoint solver

As with the aerostructural system of equations, we consider two approaches to solve the coupled adjoint system (19): a segregated approach and a monolithic approach. The segregated approach is the lagged coupled adjoint approach (Martins et al., 2004), which corresponds to a linear block Gauss–Seidel algorithm. We also investigate a monolithic approach, which uses a Krylov method applied directly to the coupled linear system.

4.5.1 Linear block Gauss–Seidel method (LBGS)

The analog of the NLBGS method for coupled adjoint equations is the linear block Gauss–Seidel (LGBS) method. This method expresses the interdisciplinary coupling as additional forcing terms to the right-hand side of each set of disciplinary adjoint equations. If we separate the single matrix equation into the two blocks we obtain:

$$\left(\frac{\partial \mathcal{A}}{\partial w} \right)^T \psi^{(k)} = \left(\frac{\partial f}{\partial w} \right)^T - \left(\frac{\partial \mathcal{S}}{\partial w} \right)^T \phi^{(k-1)} \tag{28}$$
$$\left(\frac{\partial \mathcal{S}}{\partial u} \right)_K \phi^{(k)} = \left(\frac{\partial f}{\partial u} \right)^T - \left(\frac{\partial \mathcal{A}}{\partial u} \right)^T \psi^{(k)} - \left(\frac{\partial \mathcal{S}}{\partial u} \right)_F^T \phi^{(k-1)}.$$

The subscripts K and F in the term $\partial \mathcal{S} / \partial u$ represent the contributions from the stiffness matrix and external forces, respectively. Since the external-force component is costly to compute and involves synchronous communication between disciplines, it is lagged. The same applies to the contribution from the off-diagonal term. We can iterate between these two equations until we reach the desired convergence level. The main advantage of this approach is that the discipline adjoint solvers can be re-utilized by simply adding the appropriate

right-hand side forcing terms. In practice, we only partially converge each discipline before a data exchange is made. Additionally, depending on the problem, it may be necessary to apply an under-relaxation factor to the structural update for enhanced stability. The pseudocode for this approach is listed in Algorithm 3.

Algorithm 3 Block Gauss–Seidel coupled adjoint solution method

---

1: **Given**: $\psi^{(0)}$, $\phi^{(0)}$, $k_{\max}$, $\theta \in (0,1]$ $\qquad\qquad$ ▷ $\psi^{(0)}$, $\phi^{(0)}$ can be from a previous solution

2: **if** Aerodynamic Function **then**

3: $\quad \dfrac{\partial f}{\partial u} \leftarrow T^T \left(\dfrac{\partial f}{\partial X_S} \right)$ $\qquad\qquad\qquad$ ▷ Cross-discipline function partial derivative

4: **end if**

5: **for** $k \leftarrow 1, k_{\max}$ **do**

6: $\quad \phi_A \leftarrow T\phi^{(k-1)}$ $\qquad\qquad\qquad\qquad\qquad\qquad$ ▷ Transfer structural adjoint

7: $\quad R_{\mathcal{A}}^{(k)} \leftarrow \left(\dfrac{\partial f}{\partial w} \right)^T - \left(\dfrac{\partial \mathcal{A}}{\partial w} \right)^T \psi^{(k-1)} - \left(\dfrac{\partial F_A}{\partial w} \right)^T \phi_A$ $\qquad$ ▷ Aerodynamic adjoint RHS

8: $\quad \left(\dfrac{\partial \mathcal{A}}{\partial w} \right)^T \Delta\psi = R_{\mathcal{A}}^{(k)}$ $\qquad\qquad\qquad$ ▷ Partially solve CFD adjoint update

9: $\quad \psi^{(k)} \leftarrow \psi^{(k-1)} + \Delta\psi$ $\qquad\qquad\qquad\qquad$ ▷ Update aerodynamic adjoint

10: $\quad P \leftarrow T^T \left(\dfrac{\partial X_V}{\partial X_S} \right)^T \left(\dfrac{\partial \mathcal{A}}{\partial X_V} \right)^T \psi^{(k)}$ ▷ Transfer aerodynamic contribution to structural RHS

11: $\quad Q \leftarrow T^T \left(\dfrac{\partial F_A}{\partial X_S} \right)^T T\phi^{(k-1)}$ $\qquad\qquad\qquad\qquad$ ▷ Force $\dfrac{\partial \mathcal{S}}{\partial u}$ contribution

12: $\quad R_S^{(k)} \leftarrow \left(\dfrac{\partial f}{\partial u} \right)^T - K\phi^{(k-1)} - P - Q$ $\qquad\qquad$ ▷ Compute structural RHS

13: $\quad K\Delta\phi = R_S^{(k)}$ $\qquad\qquad\qquad\qquad\qquad\qquad$ ▷ Partially solve CSM adjoint update

14: $\quad \phi^{(k)} \leftarrow \phi^{(k-1)} + \theta\Delta\phi$ $\qquad\qquad\qquad$ ▷ Under-relaxed structural adjoint update

15: $\quad$ **if** $\left\| R_{\mathcal{A}}^{(k)} \right\| < \epsilon_{SA} \left\| R_{\mathcal{A}}^{(1)} \right\|$ **and** $\| R_{\mathcal{S}}^{(k)} \| < \epsilon_{SA} \| R_{\mathcal{S}}^{(1)} \|$ **then** $\qquad$ ▷ Coupled adjoint convergence check

16: $\quad\quad$ **break**

17: $\quad$ **end if**

18: **end for**

4.5.2 Coupled Krylov method (CK)

The second approach that we use to solve the coupled adjoint equation (19) is a fully coupled monolithic method. This is the first use of a monolithic solution method for the coupled adjoint equations on a large-scale aerostructural problem.

Krylov subspace methods are particularly attractive for this type of problem, since they require only matrix-vector products, allowing the use of our pre-existing matrix-free adjoint operators. In our case, the diagonal blocks of the Jacobian in Eq. (19) are stored, but the off-diagonal terms are not explicitly stored. Effective preconditioning is critical to the performance of Krylov methods, especially on large systems of equations such as the ones we are solving.

We choose to use a block-Jacobi preconditioner for the coupled system, since this approach ignores the off-diagonal terms and allows the aerodynamic and structural preconditioning to be carried out in parallel. For the aerodynamic block preconditioner we reuse the preconditioned Krylov subspace method used for the aerodynamic adjoint. In this case however, we run a fixed number of GMRES iterations, typically between 10 and 20.

For the structural block of the preconditioner, we use the matrix factorization of K. The pseudocode for the linear adjoint operator for the coupled system is listed in Algorithm 4. The matrix-vector products are computed in a matrix-free fashion.

Algorithm 4 Coupled Krylov method linear operator

1: **function** MULT(X)   ▷ Compute Jacobian-vector product with X

2:   $(X_\mathcal{A}, X_\mathcal{S}) \leftarrow X$ ▷ Extract aerodynamic and structural components

3:   $Y_\mathcal{A} \leftarrow \left(\dfrac{\partial \mathcal{A}}{\partial w}\right)^T X_\mathcal{A}$

   $Y_\mathcal{S} \leftarrow K X_\mathcal{S}$ ▷ Evaluate diagonal contributions in parallel

4:   $Y_\mathcal{A} \leftarrow Y_\mathcal{A} + \left(\dfrac{\partial F_A}{\partial w}\right)^T T X_\mathcal{S}$ ▷ Add aerodynamic off-diagonal term

5:   $Y_\mathcal{S} \leftarrow Y_\mathcal{S} + T^T \left(\dfrac{\partial X_V}{\partial X_S}\right)^T \left(\dfrac{\partial \mathcal{A}}{\partial X_V}\right)^T X_\mathcal{A}$ ▷ Add structural off-diagonal term

6:   $Y_\mathcal{S} \leftarrow Y_\mathcal{S} + T^T \left(\dfrac{\partial F_A}{\partial X_S}\right)^T T X_\mathcal{S}$ ▷ Add force $\dfrac{\partial \mathcal{S}}{\partial u}$ contribution

7:   $Y \leftarrow (Y_\mathcal{A}, Y_\mathcal{S})$ ▷ Combine aerodynamic and structural components

8: **return** Y

9: **end function**

---

5 Aerostructural optimization applications

5.1 Multipoint wing design

Similarly to the cases presented in the previous section, we use the CRM as our baseline aircraft configuration. Since CRM geometry only provides the 1 g wing shape, we had to develop a structural wing box model and a jig wing shape, which we call the undeformed CRM (uCRM) (Kenway et al., 2014a), shown in Figure 16.

The objective for each optimization is to minimize the following function, which involves both fuel burn and takeoff gross weight (TOGW):

$$\text{Obj} = \sum_{i}^{N} \mathcal{W}_i \left(\beta \times \text{Fuel Burn} + (1.0 - \beta)\,\text{TOGW}\right). \tag{29}$$

Extreme β values of 0 and 1 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 considered in our multipoint optimization, we compute the fuel burn and TOGW as if the entire mission was completed in a cruise-climb fashion at the L/D
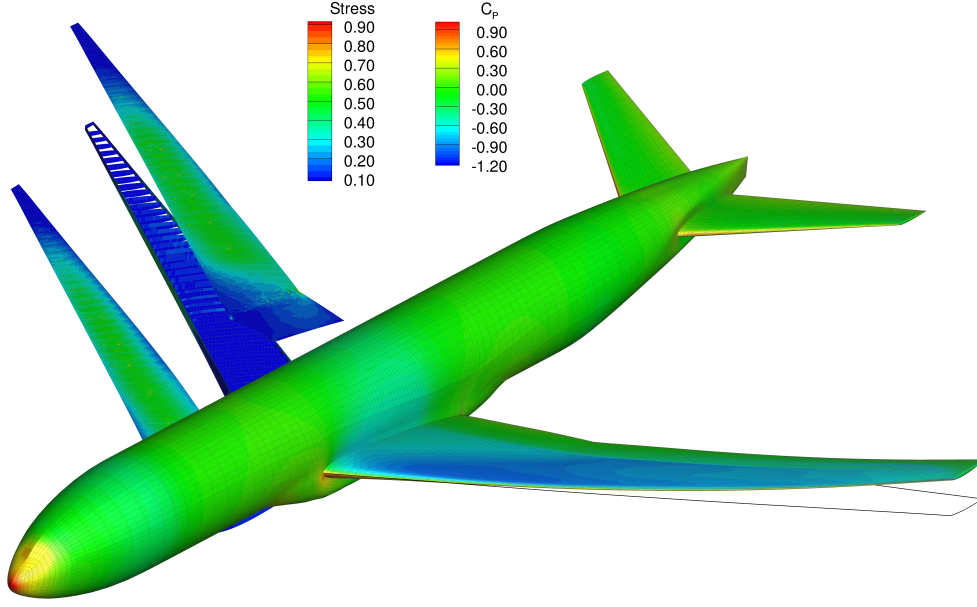
Figure 16: uCRM aerostructural model. The wingbox structural model is shown on the left and the aerodynamic model on the right. The black outline denotes the jig shape. The CFD surface mesh resolution is also visible.

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 full design range we compute the TOGW and fuel burn.

The flight conditions considered in our optimization are shown in Fig. 17. The basic operating condition is $M = 0.85$, $C_L = 0.5$, at $34\,000$ 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). We 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.

The final two operating conditions are for enforcing the buffet onset constraints, whose formulation we introduced in Sec. 3.2. 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.

In addition to the cruise and buffet operating conditions, we must also define several maneuver loading conditions that produce representative loads to size the wing structure. The three maneuver load conditions are as follows:
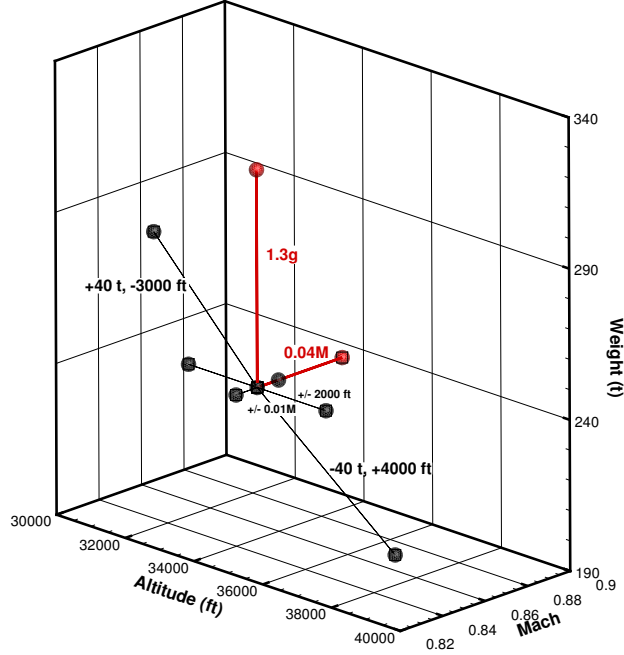
Figure 17: Flight operating conditions considered in the optimizations.

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 : 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.

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.

The full optimization problem (objective, design variables, and constraints) is summarized in Table 5.1. When it comes to design variables, there are several groups: geometric design variables, aerodynamic variables, engine model variables, and structural variables. The geometric and structural design variables are shown in Fig. 18.

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
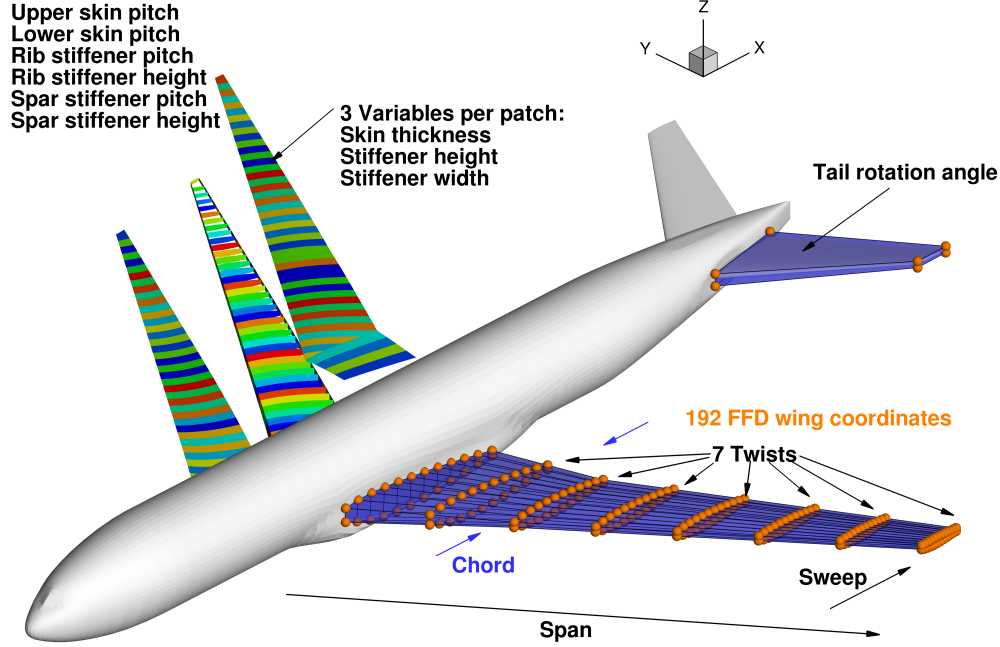
Figure 18: Geometric and structural optimization variables.

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. 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 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.

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

	Function/variable	Description	Quantity
minimize	$\sum\limits_{i}^{N} \mathcal{W}_i \left(\beta \times \text{Fuel Burn} + (1.0 - \beta)\,\text{TOGW}\right)$		
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**

Table 3: Summary of aerostructural design optimization problem.

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.

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 $-1.0\,\mathrm{g}$ 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.

We present two optimizations: one that minimizes the fuel burn ($\beta = 1$), and another that minimizes TOGW ($\beta = 0$). This enables us to see the effect of the objective choice on the optimized designs. The resulting optimizations are shown in Figs. 19 and 20.

The fuel burn minimization shown in Fig. 19 dramatically increased the wing span, at the cost of a heavier wing structure. The wingbox thickness distribution in the upper right of Fig 19 confirms that the wing skins are much thicker. There is also an increase in area, and lower thickness to chord ratios.

The TOGW minimization (Fig. 20) 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.

5.2 Tow-steered composite optimization

In the last 30 years since their introduction into aerospace applications, composites have become increasingly used, making up as much as 50% of modern aircraft by weight. Considering this fact, it is surprising that most aircraft today are only scratching the surface of the true potential of composite technology with traditional uniaxial fibers. With the introduction of automated fiber placement machines, the tow direction in laminae is now allowed to be steered spatially throughout each layer. This process is known as composite tow steering and has been shown to have improved performance over its uniaxial fiber counterpart with no additional weight penalty. With modern aircraft wings moving toward higher aspect ratios, which inevitably leads to larger deflections, it is reasonable to assume that a tow-steered composite structure can be tailored to outperform its unsteered counterpart. However, given the highly coupled nature of the aerodynamics and structural response of the problem it is not obvious nor intuitive to find the composite fiber pattern that would yield an optimum result. To address this issue, we perform the simultaneous design of aerodynamic shape, structural sizing, and tow steering angles, while considering the wing flexibility.

Two aerostructural optimizations were performed: one with the full tow steering parametrization , which we will call the *steered* case, and one where the tow orientation was held fixed, the *unsteered* case, but the laminate thickness and other variables were still allowed to vary. The unsteered case is done in order to isolate the benefits due to the tow steering process. Each optimization requires four aerostructural analysis: a cruise condition for evaluating the performance, and three conditions for which the structural constraints are enforced: a $-1.0\,\mathrm{g}$ dive, $2.5\,\mathrm{g}$ pull-up, and $1.0\,\mathrm{g}$ gust maneuver condition. The parameters for these four conditions are listed in Table 4, where TOGW is the Takeoff Gross Weight of the aircraft.
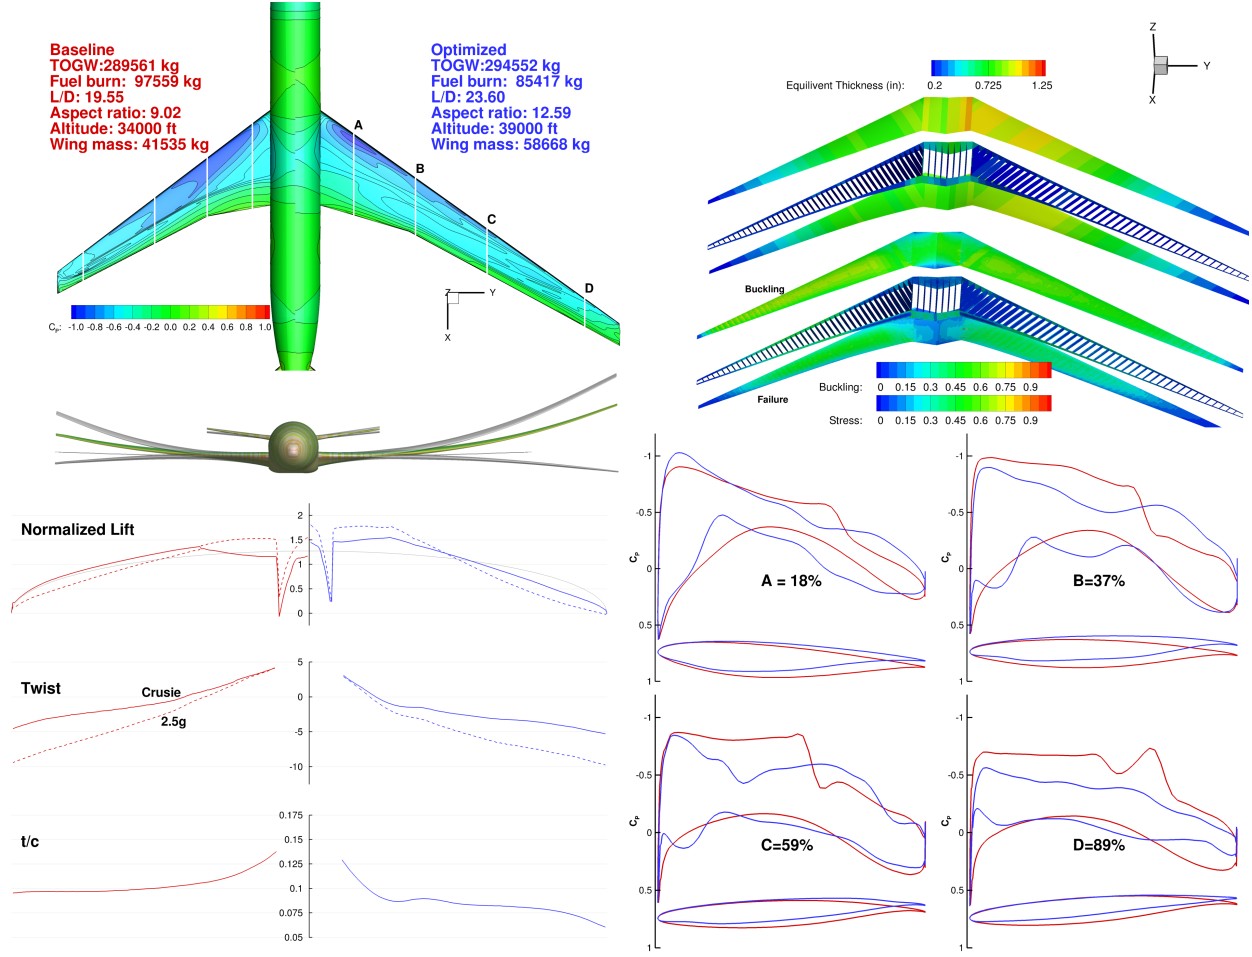
Figure 19: Fuel burn minimization result.

A summary of the tow-steered optimization problem is shown in Table 5.

Table 4: Optimization flight condition parameters

Parameter	Cruise	2.5g	-1.0g	1.0g gust
Mach number	0.85	0.64	0.64	0.86
Altitude (ft.)	37 000	0	0	27 300
Weight	LGW + 0.2 × FB	TOGW	TOGW	TOGW

As a baseline for our optimizations, we use the undeflected Common Research Model 13.5 (uCRM 13.5) is a 13.5 aspect ratio variant of the CRM. The purpose of the uCRM-13.5, similarly to that of the uCRM-9 developed by Kenway et al. (2014a), is to define a jig wing geometry and its corresponding structural layout to be used as a baseline for future aerostructural studies. The high aspect ratio wing planform was defined by taking the planform of the uCRM 9 and increasing the span while keeping the reference area constant
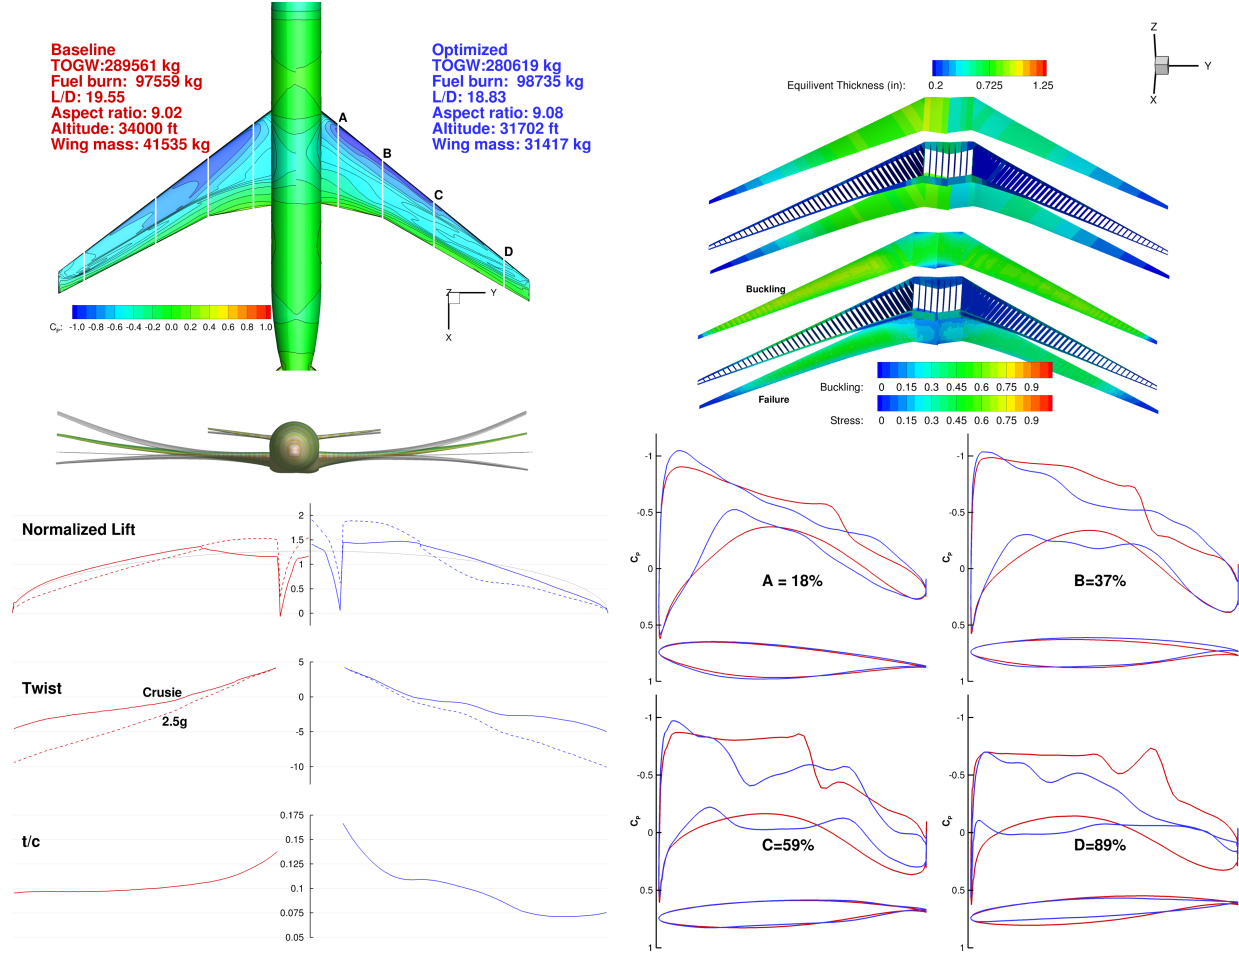
Figure 20: TOGW minimization result.

until an aspect ratio of 13.5 was achieved. The outer mold line (OML) was then defined through a single point aerostructural shape optimization of a conventional aluminum wing with the objective of minimizing fuel burn. A three-view of the uCRM-13.5 is shown in Fig. 21. The wing planform and wing box structure used for the optimization are shown in Fig. 22. In addition to aerodynamic loads the structural model also considers the inertial effects of the non-structural masses. This includes a 7,500 kg engine along with a number of leading edge and trailing edge masses concentrated at discrete locations along the span of the wing, as shown in Fig. 22. Finally, the analysis also includes the distributed weight of the fuel in the wing during flight. This inertial load is applied as a distributed traction over the lower surface of the wing skin.

The objective of the optimization problem is to minimize the fuel burn of the aircraft, calculated using the Breguet range equation.

The design variables used in the optimization can be broken down into three categories: aerodynamic, geometric, and structural. The aerodynamic design variables consists of 4 angles of attack, one for each flight condition. The geometric design variables consist of 216 wing FFD variables which control the wing cross-sectional shape, 8 twist variables,

	Variable/function	Description	Quantity
Minimize	FB	Fuel burn	
with respect to	α_i	Angle of attack for each flight case	4
	γ	Wing twist	10
	x_{shape}	Wing FFD variables	216
	ϕ_i	Tail trim angle for each case	4
	$\theta^{cp}_{\text{wing,U}}$	Wing upper skin tow variables	120
	$t^{cp}_{\text{wing,U}}$	Wing upper skin thickness variables	120
	$\theta^{cp}_{\text{wing,L}}$	Wing lower skin tow variables	120
	$t^{cp}_{\text{wing,L}}$	Wing lower skin thickness variables	120
	t_{rib}	Rib thickness variables	58
	t^{cp}_{LS}	Leading edge spar thickness variables	20
	t^{cp}_{TS}	Trailing edge spar thickness variables	20
		Total design variables	**812**
subject to	$L = n_i W$	Load factor	4
	$c^i_{m_y} = 0$	Trimmed flight	4
	$t_{LE}/t_{LE_{init}} \geq 1.0$	Leading edge radius	20
	$t_{TE}/t_{TE_{init}} \geq 1.0$	Trailing edge thickness	20
	$\text{KS}_{\text{stress}} < 1.0$	Yield stress	9
	$\text{KS}_{\text{buckling}} < 1.0$	Buckling	9
	$t > t_{\text{gage}}$	Gage thickness	258
	$-\dfrac{1}{r_{\min}} < \text{KS}_\kappa < \dfrac{1}{r_{\min}}$	Fiber path turning radius	112
	$-\psi_0 < \text{KS}_\psi < \psi_0$	Fiber path divergence	112
	$\text{KS}_{||\nabla t||} < ||\nabla t_0||$	Thickness variation	221
		Total constraints	**769**

Table 5: Tow-steered wing box minimization problem

and 4 horizontal tail incidence angles used to trim each flight condition. The structural design variables consist of 58 rib thickness variables (one for each rib) and 40 spar thickness variables (20 for each spar). The remaining structural variables for the steered case are the tow-steering variables for the upper and lower skins of the wing. These are the thickness and main tow-steering orientation values set at the control points on the upper and lower wing box skins. There are 120 control points for each surface, each of which corresponding to a thickness and tow orientation variable for a total of 480 design variables. For the unsteered optimization, the steer orientation variables are fixed while the thickness variables are free to vary. The total number of design variables for the tow-steered optimization is 812 variables.

To produce meaningful results, a number of constraints had to be applied to the optimization, as summarized in Table 5. The structural sizing of the wing box is dictated by three maneuver conditions. The first set of constraints ensure that the lift produced by the aircraft in each condition of flight had to match the weight of the aircraft in that condition multiplied by the load factor of that condition. In this case, there four lift constraints, one for each condition. In addition to this constraint, the pitching moment for each flight condition is constrained to be zero such that each flight conditioned is trimmed, for another four constraints. The leading edge radius was constrained from decreasing in order to maintain high lift performance of the wing, and the trailing edge thickness was constrained to ensure manufacturability, adding another 40 constraints. The failure and buckling constraints were aggregated into three groups: ribs and spars, upper skin, and lower skin. These constraints
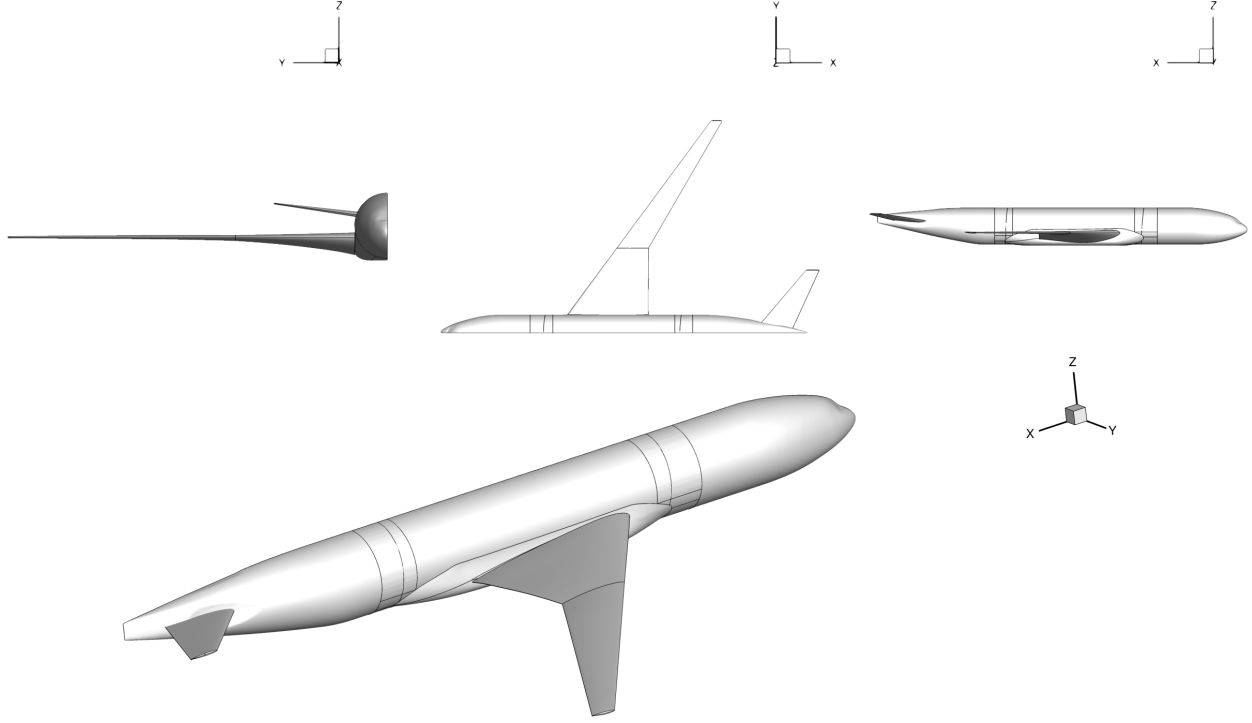
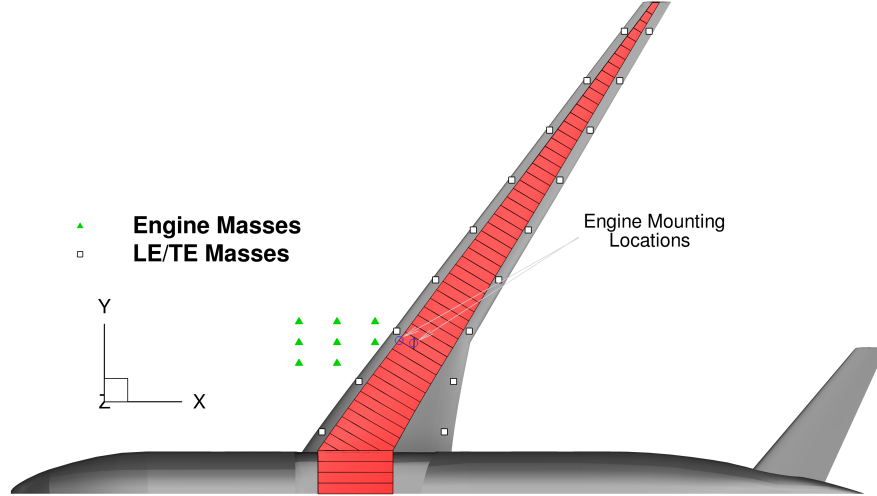Figure 21: Three-view of the uCRM-13.5 model

Figure 22: uCRM-13.5 planform and wing box with non-structural masses

were enforced for the 3 maneuver flight conditions, resulting in a total of 18 constraints: 9 for failure and 9 for buckling.

The next set of constraints are related to the divergence and curvature of the tow-steered skins. The fiber curvature was constrained to have a minimum turning radius of 70 in, which corresponds to a magnitude of curvature of $0.0143\,\mathrm{in}^{-1}$. Both constraints were aggregated over each panel of the wing box for a total of 224 constraints (112 on each side). The total

number of constraints for the optimization problem adds up to 769.

The results of the tow-steered and unsteered optimization are shown in Fig. 23. We can see that there is a 1% improvement, in the fuel burn of the tow-steered relative to the unsteered wing. From the spanwise lift distribution of the two wings for the cruise and 2.5 g maneuver, we can see that the difference between the cruise and 2.5 g maneuver conditions, which corresponds to the passive load alleviation that we seek. Although there is some amount of passive load alleviation, the steered case does not show significantly more load alleviation when compared to the unsteered case. This may be due to the fact that without stiffeners defined in the model, the buckling constraint dominates the optimization. Further evidence for this can be seen in the fact that the thickness distribution changes most noticeably on the lower skin of the wing, where the buckling would be less critical for the design. In addition, a greater difference between optimal cruise and maneuver loads is achieved for other objectives that emphasize structural weight, such as takeoff gross weight (Kenway et al., 2014a). Despite the small difference in load alleviation the optimizer is able to decrease the wing weight of the steered wing by 13% relative to the unsteered case. This is most probably because the optimizer has the freedom to place the composite tow orientation closer to the orientation of the principal stress in high stressed regions.

Figure 24 shows the a comparison between the tow-steered and unsteered patterns for each ply. Due to the buckling constraints and the absence of stringers the wing box panels are much thicker than in a realistic wing box. For both wings, the optimizer adds additional laminate thickness at the engine mounting location to deal with the local stress concentration. The main tow pattern for the upper and lower skin of the wing box both steer back towards the trailing edge at the root and the gradually sweep forward towards the tip.

6   Summary

In Part 2 of this lecture, we presented a framework for aerodynamic and aerostructural design optimization of aircraft wings. The methods used in this framework successfully tackle the compounding challenges of modeling the wing with high fidelity, while optimizing it with respect to hundreds of design variables. The effectiveness of this framework hinges the use of high-performance parallel computing, fast solvers, state-of-the-art gradient-based optimization, and an efficient and accurate approach for computing the derivatives for the aerostructural solver via the coupled adjoint method. We demonstrated the effectiveness of the methods by presenting the results of four design optimization cases. Other cases have been solved, and the interested reader can refer to other published work (Chen et al., 2016; Garg et al., 2015; James et al., 2014; Kenway and Martins, 2016b; Liem et al., 2015; Lyu and Martins, 2014, 2015; Mader and Martins, 2013).

References

Albring, T., Sagebaum, M., and Gauger, N. R. (2016). *New Results in Numerical and Experimental Fluid Mechanics X: Contributions to the 19th STAB/DGLR Symposium Munich, Germany, 2014*, chapter A Consistent and Robust Discrete Adjoint Solver for the SU2 Framework—Validation and Application, pages 77–86. Springer International Publishing, Cham.
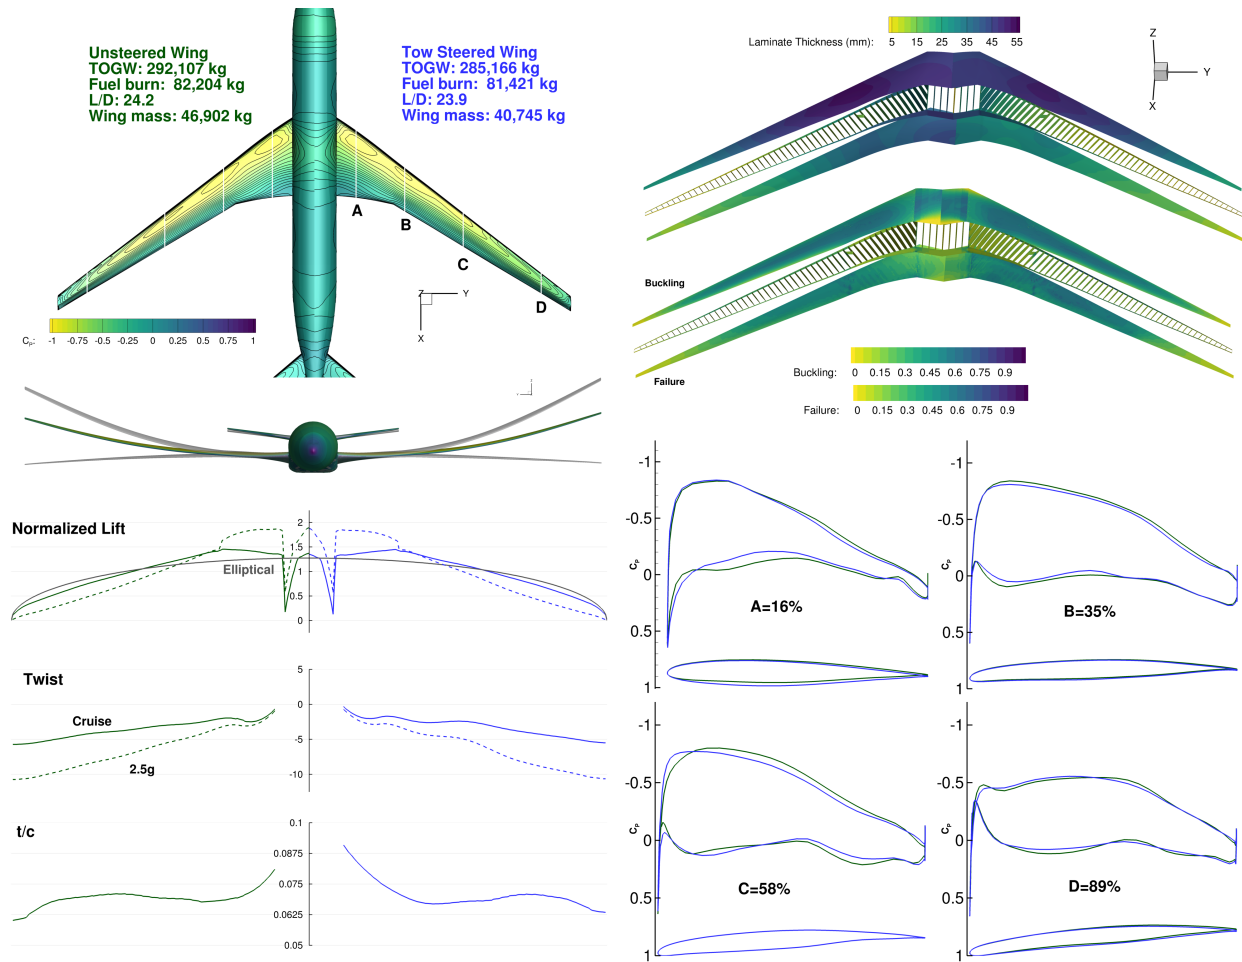
Figure 23: Comparison tow-steered (right) vs unsteered (left) aerostructural optimization.

Balay, S., Gropp, W. D., McInnes, L. C., and Smith, B. F. (1997). *Efficient Management of Parallelism in Object Oriented Numerical Software Libraries*, pages 163–202. Birkhäuser Press.

Brown, S. A. (1997). Displacement extrapolation for CFD+CSM aeroelastic analysis. In *Proceedings of the 35th AIAA Aerospace Sciences Meeting*, Reno, NV. AIAA 1997-1090.

Chen, S., Lyu, Z., Kenway, G. K. W., and Martins, J. R. R. A. (2016). Aerodynamic shape optimization of the Common Research Model wing-body-tail configuration. *Journal of Aircraft*, 53(1):276–293. doi:10.2514/1.C033328.

Garg, N., Kenway, G. K. W., Lyu, Z., Martins, J. R. R. A., and Young, Y. L. (2015). High-fidelity hydrodynamic shape optimization of a 3-D hydrofoil. *Journal of Ship Research*, 59(4):209–226. doi:10.5957/JOSR.59.4.150046.

Gill, P. E., Murray, W., and Saunders, M. A. (2002). SNOPT: An SQP algorithm for large-scale constrained optimization. *SIAM Journal of Optimization*, 12(4):979–1006. doi:10.1137/S1052623499350013.
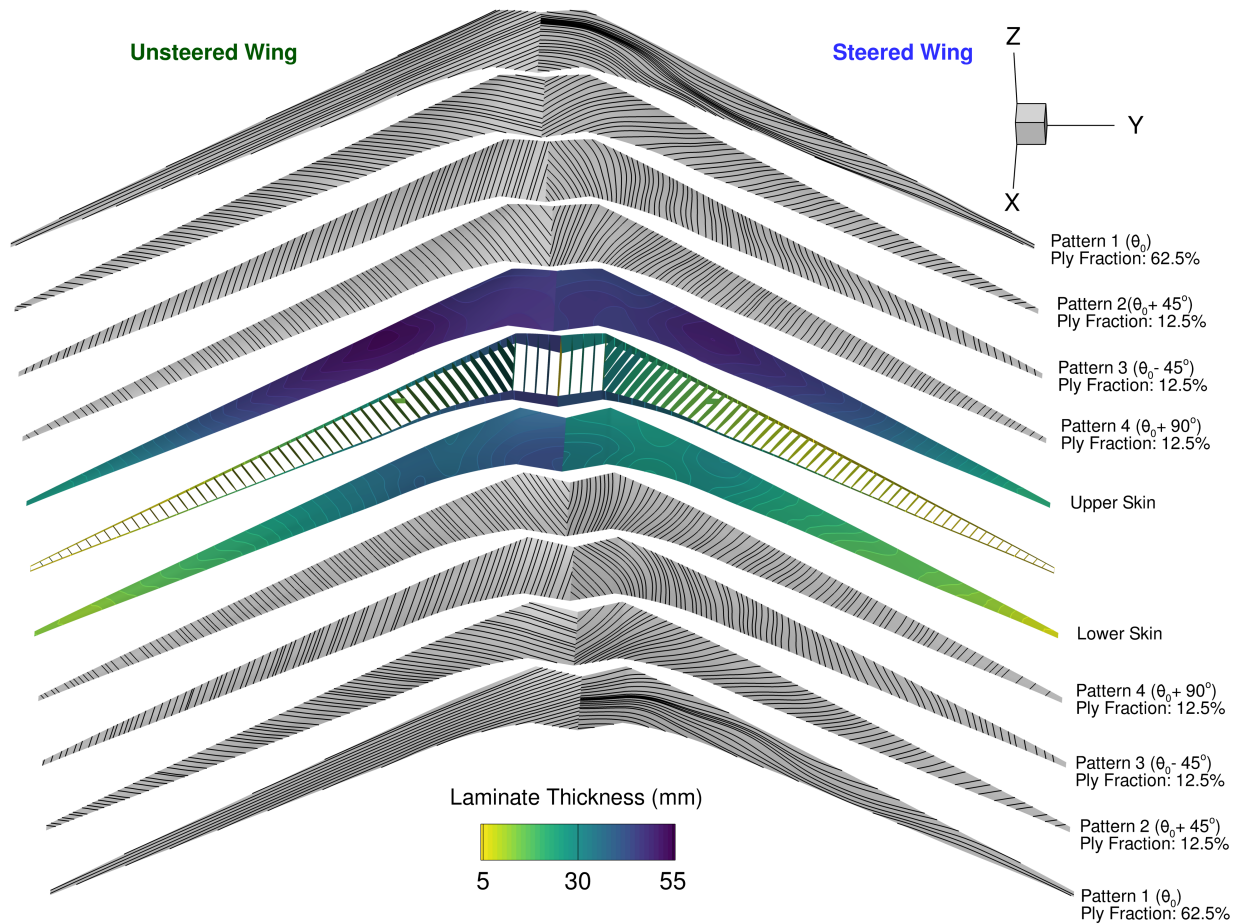
Figure 24: Comparison of optimal tow-steered vs unsteered aerostructural designs.

Haftka, R. T. (1977). Optimization of flexible wing structures subject to strength and induced drag constraints. *AIAA Journal*, 14(8):1106–1977. doi:10.2514/3.7400.

Irons, B. M. and Tuck, R. C. (1969). A version of the aitken accelerator for computer iteration. *International Journal for Numerical Methods in Engineering*, 1(3):275–277.

Irving, C. (1993). *Wide-Body—The Triumph of the 747*. William Morrow and Company, New York.

James, K. A., Kennedy, G. J., and Martins, J. R. R. A. (2014). Concurrent aerostructural topology optimization of a wing box. *Computers and Structures*, 134:1–17. doi:10.1016/j.compstruc.2013.12.007.

Jameson, A., Schmidt, W., and Turkel, E. (1981). Numerical solution of the Euler equations by finite volume methods using Runge-Kutta time stepping schemes. *AIAA Paper* 81-1259.

Kennedy, G. J. and Martins, J. R. R. A. (2014). A parallel finite-element framework for large-scale gradient-based design optimization of high-performance structures. *Finite Elements in Analysis and Design*, 87:56–73. doi:10.1016/j.finel.2014.04.011.

Kenway, G. K., Kennedy, G. J., and Martins, J. R. R. A. (2010). A CAD-free approach to high-fidelity aerostructural optimization. In *Proceedings of the 13th AIAA/ISSMO Multidisciplinary Analysis Optimization Conference*, Fort Worth, TX. AIAA 2010-9231.

Kenway, G. K. and Martins, J. R. R. A. (2016a). Aerodynamic shape optimization of the CRM configuration including buffet-onset conditions. In *54th AIAA Aerospace Sciences Meeting*, San Diego, CA. American Institute of Aeronautics and Astronautics.

Kenway, G. K. W., Kennedy, G. J., and Martins, J. R. R. A. (2014a). Aerostructural optimization of the Common Research Model configuration. In *15th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*, Atlanta, GA. AIAA 2014-3274.

Kenway, G. K. W., Kennedy, G. J., and Martins, J. R. R. A. (2014b). Scalable parallel approach for high-fidelity steady-state aeroelastic analysis and derivative computations. *AIAA Journal*, 52(5):935–951. doi:10.2514/1.J052255.

Kenway, G. K. W. and Martins, J. R. R. A. (2014). Multipoint high-fidelity aerostructural optimization of a transport aircraft configuration. *Journal of Aircraft*, 51(1):144–160. doi:10.2514/1.C032150.

Kenway, G. K. W. and Martins, J. R. R. A. (2015a). AIAA ADODG Case 5: CRM wing-body-tail optimization at flight reynolds number. Technical report, AIAA.

Kenway, G. K. W. and Martins, J. R. R. A. (2015b). Multipoint aerodynamic shape optimization investigations of the Common Research Model wing. In *Proceedings of the AIAA Science and Technology Forum and Exposition (SciTech)*, Kissimmee, FL.

Kenway, G. K. W. and Martins, J. R. R. A. (2016b). Multipoint aerodynamic shape optimization investigations of the Common Research Model wing. *AIAA Journal*, 54(1):113–128. doi:10.2514/1.J054154.

Liem, R., Kenway, G. K. W., and Martins, J. R. R. A. (2015). Multimission aircraft fuel burn minimization via multipoint aerostructural optimization. *AIAA Journal*, 53(1):104–122. doi:10.2514/1.J052940.

Lovely, D. and Haimes, R. (1999). Shock detection from computational fluid dynamics results. In *Proceedings of the 14th Computational Fluid Dynamics Conference*, Norfolk, VA.

Luke, E., Collins, E., and Blades, E. (2012). A fast mesh deformation method using explicit interpolation. *Journal of Computational Physics*, 231(2):586–601. doi:10.1016/j.jcp.2011.09.021.

Lyu, Z., Kenway, G. K., and Martins, J. R. R. A. (2015). Aerodynamic shape optimization investigations of the Common Research Model wing benchmark. *AIAA Journal*, 53(4):968–985. doi:10.2514/1.J053318.

Lyu, Z., Kenway, G. K., Paige, C., and Martins, J. R. R. A. (2013). Automatic differentiation adjoint of the Reynolds-averaged Navier–Stokes equations with a turbulence model. In *21st AIAA Computational Fluid Dynamics Conference*, San Diego, CA.

Lyu, Z. and Martins, J. R. R. A. (2014). Aerodynamic design optimization studies of a blended-wing-body aircraft. *Journal of Aircraft*, 51(5):1604–1617. doi:10.2514/1.C032491.

Lyu, Z. and Martins, J. R. R. A. (2015). Aerodynamic shape optimization of an adaptive morphing trailing edge wing. *Journal of Aircraft*, 52(6):1951–1970. doi:10.2514/1.C033116.

Mader, C. A., Kenway, G., and Martins, J. R. R. A. (2008). Towards high-fidelity aerostructural optimization using a coupled ADjoint approach. In *Proceedings of the 12th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*, Victoria, BC. AIAA 2008-5968.

Mader, C. A. and Martins, J. R. R. A. (2013). Stability-constrained aerodynamic shape optimization of flying wings. *Journal of Aircraft*, 50(5):1431–1449. doi:10.2514/1.C031956.

Mader, C. A. and Martins, J. R. R. A. (2014). Computing stability derivatives and their gradients for aerodynamic shape optimization. *AIAA Journal*, 52(11):2533–2546. doi:10.2514/1.J052922.

Martins, J. R. R. A., Alonso, J. J., and Reuther, J. J. (2004). High-fidelity aerostructural design optimization of a supersonic business jet. *Journal of Aircraft*, 41(3):523–530. doi:10.2514/1.11478.

Martins, J. R. R. A., Alonso, J. J., and Reuther, J. J. (2005). A coupled-adjoint sensitivity analysis method for high-fidelity aero-structural design. *Optimization and Engineering*, 6(1):33–62. doi:10.1023/B:OPTE.0000048536.47956.62.

Martins, J. R. R. A. and Lambe, A. B. (2013). Multidisciplinary design optimization: A survey of architectures. *AIAA Journal*, 51(9):2049–2075. doi:10.2514/1.J051895.

Martins, J. R. R. A., Mader, C. A., and Alonso, J. J. (2006). ADjoint: An approach for rapid development of discrete adjoint solvers. In *Proceedings of the 11th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*, Portsmouth, VA. AIAA 2006-7121.

Maute, K., Nikbay, M., and Farhat, C. (2001). Coupled analytical sensitivity analysis and optimization of three-dimensional nonlinear aeroelastic systems. *AIAA Journal*, 39(11):2051–2061.

Perez, R. E., Jansen, P. W., and Martins, J. R. R. A. (2012). pyOpt: A Python-based object-oriented framework for nonlinear constrained optimization. *Structural and Multidisciplinary Optimization*, 45(1):101–118. doi:10.1007/s00158-011-0666-3.

Peter, J. E. V. and Dwight, R. P. (2010). Numerical sensitivity analysis for aerodynamic optimization: A survey of approaches. *Computers and Fluids*, 39(3):373–391. doi:10.1016/j.compfluid.2009.09.013.

Poon, N. M. K. and Martins, J. R. R. A. (2007). An adaptive approach to constraint aggregation using adjoint sensitivity analysis. *Structural and Multidisciplinary Optimization*, 34(1):61–73. doi:10.1007/s00158-006-0061-7.

Saad, Y. and Schultz, M. H. (1986). GMRES: A generalized minimal residual algorithm for solving nonsymmetric linear systems. *SIAM Journal on Scientific and Statistical Computing*, 7(3):856–869. doi:10.1137/0907058.

Sederberg, T. W. and Parry, S. R. (1986). Free-form deformation of solid geometric models. *SIGGRAPH Comput. Graph.*, 20(4):151–160. doi:10.1145/15886.15903.

Spalart, P. and Allmaras, S. (1992). A one-equation turbulence model for aerodynamic flows. In *30th Aerospace Sciences Meeting and Exhibit*.

van der Weide, E., Kalitzin, G., Schluter, J., and Alonso, J. J. (2006). Unsteady turbomachinery computations using massively parallel platforms. In *Proceedings of the 44th AIAA Aerospace Sciences Meeting and Exhibit*, Reno, NV. AIAA 2006-0421.

Vassberg, J. (2011). A unified baseline grid about the common research model wing/body for the fifth aiaa cfd drag prediction workshop (invited). In *29th AIAA Applied Aerodynamics Conference*.

Vassberg, J. C., DeHaan, M. A., Rivers, S. M., and Wahls, R. A. (2008). Development of a common research model for applied CFD validation studies. AIAA 2008-6919.

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. (2014). Summary of the fourth AIAA computational fluid dynamics drag prediction workshop. *Journal of Aircraft*, 51(4):1070–1089. doi:10.2514/1.c032418.

Wrenn, G. A. (1989). An Indirect Method for Numerical Optimization Using the Kreisselmeier-Steinhauser Function. Technical report, NASA Langley Research Center, Hampton, VA.