

A Scalable Parallel Approach for High-Fidelity Aerostructural Analysis and Optimization

Gaetan K.W. Kenway,*

Graeme J. Kennedy,†

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

Joaquim R. R. A. Martins‡

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

In this paper we present several significant improvements to both coupled solution methods and sensitivity analysis techniques for high-fidelity aerostructural systems. We consider the analysis of full aircraft configurations using Euler CFD models with more than 80 million state variables and structural finite-element models with more than 1 million degrees of freedom. A coupled Newton–Krylov solution method for the aerostructural system is presented that accelerates the convergence rate for highly flexible structures. A coupled adjoint technique is presented that can compute gradients with respect to thousands of design variables accurately and efficiently. The efficiency of the presented methods is assessed on a high performance parallel computing cluster for up to 544 processors. To demonstrate the effectiveness of the proposed approach and the developed framework, an aerostructural model based on the Common Research Model is optimized with respect to hundreds of variables representing the wing outer mold line and the structural sizing. Two separate problems are solved: one where fuel burn is minimized, and another where the maximum takeoff weight is minimized. Multi-point optimizations with 5 cruise conditions and 2 maneuver conditions are performed with a 2 million cell CFD mesh and 300 000 DOF structural mesh. The optima for problems with 476 design variables are obtained within 36 hours of wall time on 435 processors. The resulting optimal aircraft are discussed and analyze the aerostructural tradeoffs for each objective.

Convergence tolerance for aerostructural solution Convergence tolerance for aerostructural adjoint solution

Nomenclature			
α	Angle of attack	I	Function of interest
ϵ_{AS}	Convergence tolerance for aerostructural solution	K	Structural stiffness matrix
ϵ_A	Convergence tolerance for aerodynamic solution	L	Lift
ϵ_S	Convergence tolerance for structural solution	R	Range of aircraft
\mathcal{A}	Aerodynamic residuals	T	Displacement transfer matrix
\mathcal{R}	All residuals	T^T	Load transfer matrix
\mathcal{S}	Structural residuals	u	Structural states
\mathcal{W}	Mesh movement operator	u_A	Displacements on CFD surface mesh
ϕ	Structural adjoint vector	V	Flight speed
ψ	Aerodynamic adjoint vector	W	Weight
θ	Under-relaxation factor	w	Aerodynamic states
c	Thrust-specific fuel consumption	W_1	Initial cruise weight
C_p	Pressure coefficient	W_2	Final cruise weight
D	Drag	X_J	Jig shape of CFD surface mesh
F	Structural load vector	x_n	Design variable vector
F_A	Aerodynamic load vector	X_S	Spatial coordinates of CFD surface mesh
		X_V	Spatial coordinates of CFD volume mesh

*Ph.D. Candidate, AIAA Student Member

†Ph.D Candidate, AIAA Student Member

‡Associate Professor, AIAA Senior Member

I. Introduction

High-fidelity aerostructural design optimization is an inherently difficult and complex problem. Typically, high-fidelity methods — such as computational fluid dynamics (CFD) and computational structural mechanics (CSM) — capture the physics of the underlying processes more faithfully, with less-restrictive assumptions, than lower-fidelity methods. Since optimization tends to exploit shortcomings of the model, high-fidelity design optimization is more likely to produce more physically realistic designs. Furthermore, high-fidelity methods provide an opportunity to perform detailed design at an earlier stage in the design process, since these methods can accurately predict the effects of detailed design changes. However, optimization problems that involve detailed design inevitably require large numbers of design variables. Given the high computational cost of high-fidelity methods, and the large numbers of design variables, gradient-based optimization with an accurate and efficient sensitivity analysis method is the only tractable approach to solving such complex design problems in a reasonable amount of time. Efficient gradient-based aerostructural design optimization of high-fidelity systems serves as the motivation for both the development of fast and robust aerostructural solution methods and an accurate and efficient coupled adjoint sensitivity analysis method.

Over the past decade, several authors have developed techniques for high-fidelity aerostructural optimization with coupled sensitivity analysis methods. Maute *et al.* [1] performed an aerostructural analysis that coupled the Euler equations to a linear finite-element model of the structure. They employed a three-field formulation where the CFD mesh movement was determined by a spring analogy. The aerodynamic loads and structural displacements were transferred between disciplines using techniques developed in the earlier work of Maman and Farhat [2] and Farhat *et al.* [3]. Maute *et al.* [1] developed a nonlinear block Gauss–Seidel method with relaxation for the solution of the aerostructural system. They presented an optimization with five design variables using gradients computed with the direct method.

Martins *et al.* [4, 5, 6] developed a coupled adjoint formulation using Euler CFD and linear finite-element analysis in a two-field formulation. They showed that the cost of computing the gradient could be made nearly independent of the number of design variables [6]. The coupled adjoint methodology was successfully applied to the optimization of a supersonic business jet with 86 design variables [5].

Later, Maute *et al.* [7] continued their previous work with the direct method and presented a coupled adjoint formulation. They showed that the accuracy of the two methods is identical. In order to improve the robustness and efficiency, Barcelos *et al.* [8] developed a class of Newton–Krylov–Schur methods for solving the coupled nonlinear fluid-structure-mesh movement problem. Their approach uses an approximate Newton method for the solution of the nonlinear coupled equations. At each iteration, a Schur complement approach is used to solve the coupled linear system that results from a linearization of the residual. They found that their technique is more robust and efficient than the original Gauss–Seidel method presented by Maute *et al.* [1]. More recently, Barcelos *et al.* [9] presented an aerostructural solution technique that couples the RANS equations to a linear structural model and a mesh movement strategy.

This paper follows the work of Martins *et al.* [5, 6, 10] and addresses many of the shortcomings of the previous work. The high computational cost and poor accuracy resulting from evaluating the off-diagonal coupled adjoint terms with finite-differencing has been eliminated. This improvement allows us to use substantially larger finite-element structural models with up to one million degrees of freedom. The aim is use a complete model of aircraft structure that is sufficiently detailed to make accurate (rather than representative) stress and failure predictions. We have also made improvements to the solution of the coupled system include the use of Aitken acceleration for the nonlinear block Gauss–Seidel method and a fully coupled Newton–Krylov solution approach that can handle highly flexible aircraft structures.

For the sensitivity analysis, a major improvement is the implementation of *all* required partial derivative terms using only methods that yield results that are accurate to machine precision. Depending on the term particular partial derivative term, analytic, forward or reverse automatic differentiation (AD) [11] sensitivity analysis methods are used. In addition to providing a more accurate sensitivity analysis, careful implementation of certain coupling terms allows for a more computationally efficient coupled adjoint solution procedure. The coupled adjoint approach presented below allows us to compute the gradient of a function of interest in approximately half the time required for an aerostructural solution. With this highly accurate and efficient method for computing sensitivity information, we were able to demonstrate the design optimization of realistic, full aircraft configurations with an unprecedented level of detail and fidelity.

The remainder of this paper is organized as follows. Section II describes the details of the aerostructural framework and Section III reviews the coupled sensitivity analysis method and solution techniques. Section IV presents a comparison of the aerostructural solution and adjoint techniques, a parallel scaling study, and a comparison between two representative high-fidelity aerostructural optimization problems.

II. Aerostructural Analysis

This section outlines the components of the aerostructural analysis, which are based on CFD analysis for the aerodynamics and a finite-element analysis for the structure. In addition to describing the two disciplinary solvers, we also describe the load and displacement transfer technique, the mesh movement approach, and the two strategies used to solve the multidisciplinary system.

1. Aerodynamic Solver

One of the goals of this project is to demonstrate the effective application of multidisciplinary design optimization (MDO) to aircraft design. Most commercial aircraft operate in the transonic flight regime and thus, at the minimum, the Euler equations must be solved to accurately predict the wave drag on the aircraft. While the inclusion of viscous effects is obviously important to accurately predict the overall drag, significant insight into the multidisciplinary problem can be gained by using the Euler equations. The CFD solver used in this work is Sumb, a structured multi-block flow solver developed by van der Weide *et al.* [12]. After discretization, we aim to solve a set of n_A nonlinear equations that takes the following form:

$$\mathcal{A}(w) = 0, \quad (1)$$

where \mathcal{A} is the vector of nonlinear aerodynamic residuals, including the inviscid and artificial dissipation terms. When the solver is used in a stand-alone mode, the equations are solved with PETSc [13] using a preconditioned matrix-free Newton–Krylov approach. A suitable initial iterate that is sufficiently close to the basin of attraction is obtained using a five-stage explicit Runge–Kutta multigrid time stepping scheme.

2. Structural Solver

The structural solver used in this work is the Toolkit for the Analysis of Composite Structures (TACS) [14]. It is a parallel finite-element code written for the analysis and design optimization of composite structures. It includes both static linear and geometrically nonlinear analysis capabilities, but only linear analysis is considered in this work. TACS is able to handle all the required structural design variables including structural thicknesses and composite ply angles. Parallelism is achieved within TACS using an element-based partition of the finite-element mesh. This partition 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 functions of interest and their sensitivities. Typically, the factorization of the stiffness matrix is the single most costly operation. To parallelize the matrix factorization and back solves, TACS uses a Schur-complemented-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 extremely poorly conditioned thin shell structural problems, with condition numbers $\mathcal{O}(10^9)$, in an efficient manner. For general nonlinear analysis the n_S structural discipline residuals, \mathcal{S} , are

$$\mathcal{S}(u) = 0. \quad (2)$$

For linear analysis this system reduces to $\mathcal{S}(u) = Ku - F$, where K is the stiffness matrix and F is the load vector.

3. Load and Displacement Transfer

The load and displacement transfer scheme we use follows the work of Brown [15], later employed by Martins *et al.* [6]. In this approach, rigid links are used to extrapolate the displacements from the structural surface to the CFD surface as shown in Figure 1. 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 occurs on the aerodynamic mesh and is transmitted back through the rigid links to the structure. The two primary advantages of this scheme are that it is consistent and conservative by construction, and that it may be used to transfer loads and displacements between aerodynamic and structural meshes that are not coincident.

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

$$\begin{aligned} X_S &= X_J + u_A \\ &= X_J + u_t + r \times u_r \end{aligned} \quad (3)$$




Figure 1: Schematic of load-displacement transfer operation

where u_t and u_r are the translational and rotational components of the displacement, u , respectively. We can compactly represent the displacement transfer from u to u_A as

$$u_A = Tu = \left(\frac{\partial u_A}{\partial u} \right) u \quad (4)$$

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

$$F = T^T F_A = \left(\frac{\partial F}{\partial F_A} \right) F_A \quad (5)$$

where F_A are the forces on the CFD nodes. More details can be found in Brown [15] and Martins *et al.* [6].

4. Mesh Movement Algorithm

CFD-based aerodynamic shape optimization requires the perturbation of the mesh to consider changes in the shape design variables. However, unless the surface perturbations are of the same order as that of the mesh off-wall spacing, at least a portion of the volume mesh must be modified in order to prevent the formation of cells with negative volumes. Aerostructural analysis imposes additional mesh movement costs that are not typically seen with aerodynamic shape optimization. Whereas aerodynamic shape optimization typically only adapts the mesh once per design cycle with an updated shape, aerostructural analysis requires a mesh movement for every displacement transfer. $\mathcal{O}(10^2)$ mesh movements may be necessary during a single aerostructural analysis, depending on the solution strategy. Thus, it is extremely important to have an efficient mesh movement algorithm. We use a hybrid algebraic-linear elasticity method for the mesh movement similar to Martins *et al.* [5] and Hicken and Zingg [16]. The idea behind this approach is to apply a linear elasticity based scheme to a coarse approximation of the mesh to account for large, low frequency perturbations, and to use an algebraic warping approach to attenuate small, high frequency perturbations. A mesh perturbation proceeds by first selecting a coarse approximation of the mesh known as *super nodes* and computing the location of the interior super node unknowns using a linear elasticity method. Next, an intermediate mesh is computed using the coarse approximation with a trivariate Bézier cubic spline interpolation algorithm. Finally, any deviations between the desired surface and the intermediate interpolation approximation is attenuated using a modified transfinite interpolation algorithm operating on the difference between the two surfaces. The goal is to compute a similar, high quality perturbed mesh like the one that would be obtained by using a linear elasticity scheme for all cells, but with a much lower computational cost. The mesh movement algorithm employed in this work is described in more detail in previous work by the authors [17].

5. Aerostructural Analysis Methods

The aerostructural optimization in this paper can be viewed as a two-discipline MDO problem. There are a number of possible MDO architectures that can be used, including single level architectures such as Multidisciplinary Design Feasible (MDF) and Individual Design Feasible (IDF), as well as decomposed architectures, such as Collaborative

Optimization (CO) or Concurrent Subspace Optimization (CSSO) [18]. Due to the highly coupled nature of aerostructural analysis and the need for efficient coupled sensitivity analysis, we select the MDF approach [19]. This approach hides the multidisciplinary nature of the problem from the optimizer and we must ensure that the system is converged to a feasible state at each design iteration.

Using a two field MDO formulation, we can write the governing equations of both disciplines as function of the fluid states, w , structural states, u , and the design parameters x_n . The latter is the vector of design variables, which consists of global variables and local variables. The global variables affect the two disciplines directly and in this case they are the geometric parameters defining changes to the aircraft surface and internal structure. The local variables only directly affect a single discipline. The angle of attack, α , is an aerodynamic local variable while structural skin thicknesses are local structural variables. Combining the residual equations from the aerodynamic and structural disciplines we write the combined residual of the multidisciplinary system as

$$\mathcal{R} = \begin{bmatrix} \mathcal{A}(w, u, x_n) \\ \mathcal{S}(u, w; x_n) \end{bmatrix} = 0. \quad (6)$$

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

5.1. Nonlinear Block Gauss–Seidel Method (NLBGS)

The typical process for solving the coupled aerostructural equations (6) is to use a nonlinear block Gauss–Seidel (NLBGS) method [6, 1]. The general solution procedure is listed in Algorithm 1. Three tolerances are defined for the coupled analysis: the relative tolerances, $\epsilon_{\mathcal{A}}$ and $\epsilon_{\mathcal{S}}$ are the tolerances to which the aerodynamic and structural disciplines are converged for each NLBGS iteration; the aerodynamic solver tolerance, $\epsilon_{\mathcal{A}}$, is typically $\mathcal{O}(10^{-1})$, while $\epsilon_{\mathcal{S}}$ is typically $\mathcal{O}(10^{-3})$ or smaller. The overall aerostructural solution tolerance, ϵ_{AS} , represents the feasibility of the interdisciplinary coupling. The CFD convergence criteria is evaluated at the beginning of the k^{th} iteration by using the structural displacements and aerodynamic states from the previous iteration. The structural convergence is evaluated at the beginning of the structural solution by using the previous structural states with the current force vector. Aitken acceleration [20] (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 discipline solver can be used without modification. For tightly coupled aerostructural problems with large displacements, however, this method may converge slowly or not at all as observed by Barcelos *et al.* [8].

5.2. Newton–Krylov Method (NK)

A second approach considered for the multi-disciplinary solution the fully coupled Newton–Krylov (NK) method. Our implementation follows closely previous work by the authors on lower fidelity multidisciplinary systems [14]. Algorithm 2 gives the procedure for evaluating the coupled nonlinear residual, \mathcal{R} .

Algorithm 2 Coupled nonlinear residual

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

The procedure is similar to one iteration of the NLBGS method with the exception that only the residual is evaluated, rather than doing an approximate solution and updating the state variables. Provided the mesh movement strategy is efficient, the coupled residual evaluation is inexpensive. We choose to use an inexact Newton–Krylov approach for solving the coupled equations (6) by computing the approximate Newton update,

$$\begin{bmatrix} \frac{\partial \mathcal{A}}{\partial w} & \frac{\partial \mathcal{A}}{\partial u} \\ \frac{\partial \mathcal{S}}{\partial w} & \frac{\partial \mathcal{S}}{\partial u} \end{bmatrix} \begin{bmatrix} \Delta w \\ \Delta u \end{bmatrix} = - \begin{bmatrix} \mathcal{A}(w) \\ \mathcal{S}(u) \end{bmatrix}. \quad (7)$$

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:    $X_S^{(k)} \leftarrow T u^{(k-1)} + X_J$  ▷ Transfer displacements from CSM to CFD
5:    $X_V^{(k)} \leftarrow \mathcal{W} \left( X_S^{(k)} \right)$  ▷ Move 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  $\|\mathcal{A}(w_i; x_n)\| > \epsilon_{\mathcal{A}} \|\mathcal{A}^{(k)}\|$  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  $\|\mathcal{S}(u^{(k)}; x_n)\| > \epsilon_{\mathcal{S}} \|\mathcal{S}^{(k)}\|$  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:   $u^{(k)} \leftarrow u_i^{(k-1)}$  ▷ Temporary new CSM solution
20:   $\Delta u^{(k)} \leftarrow u^{(k)} - u^{(k-1)}$  ▷ Compute displacement increment
21:  if  $k > 1$  then
22:     $\theta \leftarrow \theta \left( 1 - \frac{(\Delta u^{(k)} - \Delta u^{(k-1)}) \cdot \Delta u^{(k)}}{\|\Delta u^{(k)} - \Delta u^{(k-1)}\|^2} \right)$  ▷ 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  $\|\mathcal{A}^{(k)}\| < \epsilon_{AS} \|\mathcal{A}^{(1)}\|$  and  $\|\mathcal{S}^{(k)}\| < \epsilon_{AS} \|\mathcal{S}^{(1)}\|$  then ▷ Aerostructural convergence criterion
26:    break
27:  end if
28: end for
```

Once this update is computed, the state variables w and u are updated with Δw and Δu , respectively. Line search strategies are often used to increase the robustness of the Newton method however, we have observed slow convergence while using a line search procedure. For this reason we accept the full Newton step and continue.

To reduce the memory requirement, a matrix-free method is used. The matrix-vector products required by the Krylov method are approximated with finite-difference Fréchet derivatives. In theory, with this approach, only residual evaluations are required to solve the coupled problem. However, Krylov methods require effective preconditioning for acceptable performance, especially with very large systems with millions of degrees of freedom. We have implemented a block-Jacobi preconditioner that reuses code from the original discipline solvers. This has two advantages: 1) the application of the preconditioners from each discipline can be applied in parallel and it eliminates contributions from the off-diagonal terms, which are never explicitly stored and 2) since both of our discipline codes use implicit methods for their own solutions, this preconditioner simply reuses the same linear solution method from the original solver. Since these preconditioners vary from iteration to iteration we must use a flexible variant of a Krylov method to solve the coupled Newton update (7). In our case, we choose to use FGMRES [21], which has been shown to work well for a wide range of large unsymmetric systems. The Eisenstat-Walker method is used to adaptively select the forcing tolerance of the linear system solution to prevent over solving during the initial iterations [22].

As with most Newton-type methods, an arbitrary initial condition may or may not converge to the actual solution. We have not investigated the start up behavior for the coupled Newton–Krylov method in detail yet. One possible approach is to run a few relaxed NLBGS iterations from Algorithm 1 before switching to the Newton–Krylov al-

gorithm. This has worked for the purposes of this research, but it is possible other methods may be more robust especially for even more flexible designs. One main advantage of the Newton–Krylov method is the elimination of the under-relaxation parameter, which leads to faster convergence for problems with large deformations.

6. Skin Friction Estimate

Since we are using the Euler equations to model flow over the aircraft, we are not able to evaluate the skin friction drag contribution. However, the aerostructural optimizations presented later in the paper require the skin friction component in order to include planform shape variables. Otherwise, the drag penalty associated with an increase in wing area and small wing chords would not be considered, which would result in unrealistic planforms. Therefore, we use a flat plate turbulent skin friction estimate with form factor corrections to account for non-zero thicknesses. The van Driest II Method [23] is used to estimate the turbulent skin friction coefficient, C_f . For wing-like bodies, we use the following form-factor correction [24],

$$FF = 1.0 + 2.7 \left(\frac{t}{c} \right) + 100 \left(\frac{t}{c} \right)^4 \quad (8)$$

where t/c is the thickness to chord ratio. For fuselage like bodies, the correction is

$$FF = 1.0 + 1.5 \left(\frac{d}{l} \right)^{1.5} + 50 \left(\frac{d}{l} \right)^3, \quad (9)$$

where d/l is the ratio of diameter to length. The contribution of a given component to the drag coefficient is then

$$C_D = C_f \frac{A_{\text{wet}}}{A_{\text{ref}}} FF \quad (10)$$

where W_{textwet} is the wetted surface area.

In our implementation, we use 3D surface geometries as shown in Figure 2, which allows us to use same geometric design variables as the aerodynamic and structural disciplines. The skin friction drag for the wing and tail computed in a strip-wise fashion that accounts for local changes in the Reynolds number due to chord modifications, and thickness to chord changes due to shape changes. Smooth maximum functions are used to estimate the chord and t/c ratio from the discrete surface data to ensure continuous derivatives.




Figure 2: The three separate meshes used for the skin friction drag estimation

7. Framework Integration

The software architecture for high-fidelity multi-disciplinary problems is inherently complex. We make extensive use of the Python programming language as high-level scripting “glue” language [10]. Most of the framework components (flow solver, mesh warping, skin friction, and optimizer) are written in Fortran and are subsequently wrapped with `f2py` [25, 10] to provide a consistent Python interface. The structural solver and load-displacement transfer operations are written in C++ and are wrapped with SWIG. The Python modularity makes it easier to easily

couple codes together, while simultaneously maintaining their own stand-alone functionality. This way of wrapping codes allows direct memory access to the underlying program from Python, eliminating any potentially costly and error prone file I/O data transfer. The use of direct memory access allows us to efficiently implement monolithic solution techniques for both the nonlinear aerostructural system and adjoint systems, while maintaining the flexibility of a modular software architecture. Furthermore, by encapsulating the underlying compiled analysis as Python objects we can setup very complex multi-point optimization problems with no changes to the underlying compiled code. The analysis tools used in this work are components of the MDO for aircraft configurations with high fidelity, (MACH) framework.

III. Coupled Adjoint Sensitivity Analysis

1. Coupled Adjoint Overview

One critical component of the MACH framework is an efficient sensitivity analysis method. Using the approach outlined by Martins *et al.* [6], we write the adjoint equations for a coupled multidisciplinary system. A more general derivation of this theory and its connection to other sensitivity analysis methods can be found in Martins and Hwang [26]. The residuals, state variables and adjoint variables for the aerodynamic discipline are \mathcal{A} , w , ψ and for the structural discipline are \mathcal{S} , u , ϕ . As with a single discipline adjoint formulation, we can write the total sensitivity equations for the coupled system. First, we write the total sensitivity of the objective,

$$\frac{dI}{dx_n} = \frac{\partial I}{\partial x_n} - \begin{bmatrix} \frac{\partial I}{\partial w} & \frac{\partial I}{\partial u} \end{bmatrix} \begin{bmatrix} \frac{dw}{dx_n} \\ \frac{du}{dx_n} \end{bmatrix}. \quad (11)$$

We make a distinction between the partial derivatives denoted by “ ∂ ” and the “total” derivatives denoted with “ d ”. The total derivatives denote the fact the derivative requires the solution of the system of equations. We write the total sensitivity of the residuals as,

$$\begin{bmatrix} \frac{d\mathcal{A}}{dx_p} \\ \frac{d\mathcal{S}}{dx_n} \end{bmatrix} = \begin{bmatrix} \frac{\partial \mathcal{A}}{\partial x_p} \\ \frac{\partial \mathcal{S}}{\partial x_n} \end{bmatrix} - \begin{bmatrix} \frac{\partial \mathcal{A}}{\partial w} & \frac{\partial \mathcal{A}}{\partial u} \\ \frac{\partial \mathcal{S}}{\partial w} & \frac{\partial \mathcal{S}}{\partial u} \end{bmatrix} \begin{bmatrix} \frac{dw}{dx_n} \\ \frac{du}{dx_n} \end{bmatrix} = 0. \quad (12)$$

Substituting the solution of Equation (12) into Equation (11) to eliminate the total derivatives we arrive at,

$$\frac{dI}{dx_n} = \frac{\partial I}{\partial x_n} - \underbrace{\begin{bmatrix} \frac{\partial I}{\partial w} & \frac{\partial I}{\partial u} \end{bmatrix} \begin{bmatrix} \frac{\partial \mathcal{A}}{\partial w} & \frac{\partial \mathcal{A}}{\partial u} \\ \frac{\partial \mathcal{S}}{\partial w} & \frac{\partial \mathcal{S}}{\partial u} \end{bmatrix}^{-1}}_{\Psi} \begin{bmatrix} \frac{\partial \mathcal{A}}{\partial x_p} \\ \frac{\partial \mathcal{S}}{\partial x_n} \end{bmatrix}. \quad (13)$$

There are two techniques for solving this. One option is to solve Equation (12) for $[dw/dx_n, du/dx_n]^T$ once for design variable x_n and use Equation (11) to compute the total sensitivity. This is known as the direct method. If, however, there are more design variables than functions of interest, which is most likely the case with high-fidelity aerostructural optimization, it is more computationally efficient to use the adjoint method and solve for the coupled adjoint vector, Ψ . Rearranging Equation (13) gives

$$\begin{bmatrix} \frac{\partial \mathcal{A}}{\partial w} & \frac{\partial \mathcal{A}}{\partial u} \\ \frac{\partial \mathcal{S}}{\partial w} & \frac{\partial \mathcal{S}}{\partial u} \end{bmatrix}^T \begin{bmatrix} \psi \\ \phi \end{bmatrix} = \begin{bmatrix} \frac{\partial I}{\partial w} \\ \frac{\partial I}{\partial u} \end{bmatrix}. \quad (14)$$

These are the coupled adjoint equations, which are solved once for each function of interest. After the solution for the coupled adjoint equations (14) is obtained, the following equation can be used to determine the total sensitivity,

$$\frac{dI}{dx_n} = \frac{\partial I}{\partial x_n} - \left(\frac{\partial \mathcal{A}}{\partial x_n} \right)^T \psi - \left(\frac{\partial \mathcal{S}}{\partial x_n} \right)^T \phi. \quad (15)$$

The implementation details for the required partial derivative terms are explained in the following sections.

2. Coupled Adjoint Implementation

An accurate and efficient solution of the coupled adjoint equations (14) that is scalable in a parallel computing environment is challenging. In this section, we present some important aspects involved in the solution of Equation (14), as well as the computation of the required partial derivative terms. We focus especially on trade-offs between cost of implementation, computational effort and memory requirements.

2.1. Aerodynamic Residual Partial Derivatives

The aerodynamic diagonal block, $\partial\mathcal{A}/\partial w$, of the coupled adjoint equations, which represents the derivative of the aerodynamic residuals with respect to the aerodynamic states, can be used in an unmodified form from the aerodynamic solver's adjoint formulation. This term is computed exactly using forward mode automatic differentiation [27]. An analytic matrix coloring scheme for the fixed structured stencils utilized reduce the number of function evaluations required to fully populate the matrix [28].

The second term in the first row of the coupled adjoint Jacobian is $\partial\mathcal{A}/\partial u$, which represents the sensitivities of the aerodynamic residual with respect to the structural displacements. This is a challenging term to compute in a 2-state aerostructural formulation since it is a large matrix, and the depending on the mesh movement strategy involved, may be either sparse or nearly dense. We compute this term using the chain rule

$$\left(\frac{\partial\mathcal{A}}{\partial u}\right)^T \psi = \left(\frac{\partial X_S}{\partial u}\right)^T \left(\frac{\partial X_V}{\partial X_S}\right)^T \left(\frac{\partial\mathcal{A}}{\partial X_V}\right)^T \psi \quad (16)$$

where the first term from the right, $\partial\mathcal{A}/\partial X_V$, contains the sensitivities of the aerodynamic residual with respect to variations of all volume mesh coordinates. This term is computed and stored using forward mode AD, in a similar fashion to $\partial\mathcal{A}/\partial w$, such that the transpose matrix vector products with this matrix are relatively inexpensive.

The second term, $\partial X_V/\partial X_S$, is the matrix of sensitivities derivative of the mesh warping procedure which requires a careful implementation to be computational efficient. For the transpose matrix vector products required for the adjoint computation, the normal warping procedure proceeds in reverse. Reverse mode AD is used to compute the variation of the all surface nodes as well as the variation of the volume super nodes. The variation of the volume super nodes is related to the surface super nodes through the solution of a linear system of equations. We then solve a mesh adjoint equation to complete the variation of the surface super nodes. This computation only needs to be performed 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 movement is symmetric and fixed, a parallel LU decomposition is performed once using SuperLU_DIST [29]. Subsequent mesh perturbations and adjoint solutions can be computed quickly with a simple back solve operation.

Finally, since X_S , the displaced aerodynamic surface, is given by $X_S = X_J + Tu$, the partial term $(\partial X_S/\partial u)^T$ is simply T^T , the standard load transfer operation. This "reverse" computation contrasts the technique used by Martins *et al.* [5], where finite-differences were used to perturb each structural degree of freedom on the surface of the finite-element model, move the mesh, and re-evaluate the aerodynamic residual. The computational cost of this method scales directly with the number of structural surface degrees of freedom. In that work, 132 structural surface nodes yielding 396 degrees of freedom were used and it became clear that the method would not scale to the approximately 96 000 surface DOF routine used in this work.

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

$$\frac{\partial\mathcal{A}}{\partial x_n} = \left(\frac{\partial X_S}{\partial x_n}\right)^T \left(\frac{\partial X_V}{\partial X_S}\right)^T \left(\frac{\partial\mathcal{R}}{\partial X_V}\right)^T \psi. \quad (17)$$

The right-most two terms are identical to the computation for $\partial\mathcal{A}/\partial u$, and the computation of remaining term, $\partial X_S/\partial x_n$, is described in Section 2.4.

2.2. Structural Residual Partial Derivatives

The off-diagonal structural residual matrix, $\partial\mathcal{S}/\partial w$, contains the sensitivities of the structural residuals with respect to the aerodynamic states. The only contribution to the structural residual is from the aerodynamic forces, F_A , and is given by

$$\begin{aligned} \left(\frac{\partial\mathcal{S}}{\partial w}\right)^T \phi &= - \left(\frac{\partial F}{\partial w}\right)^T \phi \\ &= - \left(\frac{\partial F_A}{\partial w}\right)^T \left(\frac{\partial F}{\partial F_A}\right)^T \phi \\ &= - \left(\frac{\partial F_A}{\partial w}\right)^T T\phi. \end{aligned} \quad (18)$$

We compute and store $\partial F_A/\partial w$, as well as $\partial F_A/\partial X_S$, which we refer to as the *coupling matrices*. These terms are computed using reverse mode AD on the CFD force evaluation routine.

The remaining diagonal term, $\partial\mathcal{S}/\partial u$ represents the derivative of the structural residuals with respect to the structural states. For linear structural analysis, this term is *not* simply the linear stiffness matrix, K . Given the structural residuals (2), we find that the applied forces, F , are actually explicitly 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 produce accurate derivatives this effect must be included. Given the Krylov-type solution strategy, we compute this term in the following manner:

$$\begin{aligned} \left(\frac{\partial\mathcal{S}}{\partial u}\right)^T \phi &= \left(K^T - \left(\frac{\partial F}{\partial u}\right)^T\right) \phi \\ &= \left(K^T - \left(\frac{\partial F}{\partial F_A}\right) \left(\frac{\partial F_A}{\partial X_S}\right)^T \left(\frac{\partial X_S}{\partial u}\right)\right) \phi \\ &= \left(K^T - T^T \left(\frac{\partial F_A}{\partial X_S}\right)^T T\right) \phi. \end{aligned} \quad (19)$$

The partial derivative of the structural residuals with respect to the design variables, $\partial\mathcal{S}/\partial x_n$, is evaluated analytically.

2.3. Objective Function Partial Derivatives

The right hand side of coupled adjoint system is the derivative of the objective with respect to the system states. We consider only aerodynamic or structural objectives, that is, objectives that could be used in a single discipline analysis. The sensitivities of typical aerodynamic metrics such as C_L, C_D and C_M , are computed with reverse mode AD, in a similar fashion to the coupling matrices. The derivatives of typical structural functions, such as individual element stresses, σ_i , or the Kreisselmeier–Steinhauser (KS) [30, 31, 32] function are evaluated analytically. For structural objectives, $\partial I/\partial w$ is exactly zero. However, for aerodynamic objectives involving surface pressure or traction integration, $\partial I/\partial u$ is non-zero and is given by

$$\begin{aligned} \frac{\partial I}{\partial u} &= \left(\frac{\partial X_S}{\partial u}\right)^T \left(\frac{\partial I}{\partial X_S}\right) \\ &= T^T \left(\frac{\partial I}{\partial X_S}\right). \end{aligned} \quad (20)$$

The resulting effect is identical to the $\partial\mathcal{S}/\partial u$ calculation in Section 2.2.

2.4. Total Sensitivity Partial Derivatives

The evaluation of the partial derivative terms in Equation (15) is challenging because the load and displacement transfer operation has a dependency on geometric-type design variables. Figure 1 show how the displaced surface coordinates, X_S , are extrapolated from the structural domain. The complication arises from the fact that the length of extrapolation vector, r , changes with the design variables. This leads to a transfer scheme that is more robust to large geometry changes, but complicates the sensitivity analysis. For aerodynamic objective functions, we compute $\partial I/\partial X_S$ using reverse mode AD (similar to what was done for the coupling matrices). Then, for each design variable, we evaluate the derivative of the perturbed shape X_S with respect to the design variables, x_n , by computing

$$\begin{aligned} \frac{\partial I}{\partial x_n} &= \frac{\partial X_S}{\partial x_n}^T \frac{\partial I}{\partial X_S} \\ &= \left(\frac{\partial X_J}{\partial X_n} + \frac{\partial(r \times u_r)}{\partial x_n}\right)^T \frac{\partial I}{\partial X_S}. \end{aligned} \quad (21)$$

Since the transfer operation matrix is never explicitly formed, we evaluate a single entry of the resulting vector in sequence 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 $\partial X_S/\partial x_n$ term is used to complete the $\partial\mathcal{A}/\partial x_n$ computation. These operations are relatively costly since they require the synchronous transfer of data from both disciplines.

3. Coupled Adjoint Solution Strategy

As with the multi-disciplinary analysis system of equations, we can choose two types of solution techniques to solve the coupled sensitivity analysis system (14): a segregated approach or a monolithic approach. The most common segregated approach is the linear block Gauss–Seidel method, also known as the lagged coupled adjoint approach [5]. We also investigate a monolithic approach which uses a Krylov method applied directly to the coupled linear system.

3.1. Linear Block Gauss–Seidel Method (LBGS)

The Gauss–Seidel method expresses the interdisciplinary coupling as additional forcing terms to the right hand side of each respective adjoint solver. If we separate the single matrix equation into the two block equations we obtain the following:

$$\begin{aligned} \left(\frac{\partial \mathcal{A}}{\partial w}\right)^T \psi^{(k)} &= \frac{\partial I}{\partial w} - \left(\frac{\partial \mathcal{S}}{\partial w}\right)^T \phi^{(k-1)} \\ \left(\frac{\partial \mathcal{S}}{\partial u}\right)_K^T \phi^{(k)} &= \frac{\partial I}{\partial u} - \left(\frac{\partial \mathcal{A}}{\partial u}\right)^T \psi^{(k-1)} - \left(\frac{\partial \mathcal{S}}{\partial u}\right)_F^T \phi^{(k-1)} \end{aligned} \quad (22)$$

The subscripts K and F on the partial 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 a synchronous communication between disciplines, it is lagged, as is the contribution from the off-diagonal term. We can simply 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 term. In practice, we only partially converged each discipline before a data exchange is made. Additionally, depending on the problem, it may be necessary to apply a relaxation factor to the structural update. A description this approach is given in Algorithm 3.

Algorithm 3 Block Gauss–Seidel coupled adjoint solution method

- 1: Given: $\psi^{(0)}, \phi^{(0)}, k_{\max}$ ▷ $\psi^{(0)}$ and $\phi^{(0)}$ are possibly initialized from previous solution
 - 2: Set: $\theta \in (0, 1]$
 - 3: **if** Aerodynamic Objective **then**
 - 4: $\frac{\partial I}{\partial u} \leftarrow T^T \left(\frac{\partial I}{\partial X_S} \right)$ ▷ Cross-discipline objective partial
 - 5: **end if**
 - 6: **for** $k \leftarrow 1, k_{\max}$ **do**
 - 7: $\phi_A \leftarrow T \phi^{(k-1)}$ ▷ Transfer structural adjoint
 - 8: $R_A^{(k)} \leftarrow \frac{\partial I}{\partial w} - \left(\frac{\partial \mathcal{A}}{\partial w}\right)^T \psi^{(k-1)} - \left(\frac{\partial F_A}{\partial w}\right)^T \phi_A$ ▷ Aerodynamic adjoint RHS
 - 9: $\left(\frac{\partial \mathcal{A}}{\partial w}\right)^T \Delta \psi = R_A^{(k)}$ ▷ Partially solve CFD adjoint update
 - 10: $\psi^{(k)} \leftarrow \psi^{(k-1)} + \Delta \psi$ ▷ Update aerodynamic adjoint
 - 11: $P \leftarrow T^T \left(\frac{\partial X_V}{\partial X_S} \right)^T \left(\frac{\partial \mathcal{A}}{\partial X_V} \right)^T \psi^{(k)}$ ▷ Transfer aerodynamic contribution to structural RHS
 - 12: $Q \leftarrow T^T \left(\frac{\partial F_A}{\partial X_S} \right)^T T \phi^{(k-1)}$ ▷ Force $\frac{\partial \mathcal{S}}{\partial u}$ contribution
 - 13: $R_S^{(k)} \leftarrow \frac{\partial I}{\partial u} - K \phi^{(k-1)} - P - Q$ ▷ Compute Structural RHS
 - 14: $K \Delta \phi = R_S^{(k)}$ ▷ Partially solve CSM adjoint update
 - 15: $\phi^{(k)} \leftarrow \phi^{(k-1)} + \theta \Delta \phi$ ▷ Under-relaxed structural adjoint update
 - 16: **if** $\|R_A^{(k)}\| < \epsilon_{SA} \|R_A^{(1)}\|$ **and** $\|R_S^{(k)}\| < \epsilon_{SA} \|R_S^{(1)}\|$ **then** ▷ Coupled adjoint convergence check
 - 17: Break
 - 18: **end if**
 - 19: **end for**
-

3.2. Coupled Krylov Method (CKM)

The second type of approach that can be used to solve the coupled sensitivity equation (14) is to use a fully coupled, monolithic method. As with the Newton–Krylov solution method, this method follows previous work with lower

fidelity systems [14]. We believe this is the first use of a monolithic solution method for the coupled adjoint equations on large-scale, aerostructural problems.

Krylov subspace methods, are particularly attractive for this type of problem, as they only require matrix-vector products, allowing the use of our pre-existing matrix-free adjoint operators. In our case, the diagonal blocks of the Jacobian in Equation (14) 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 as 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, only the application of K^{-1} is used. The specification of the linear adjoint operator for the coupled system is described in in Algorithm 4. Matrix-vector products are computed, as required, for these terms in a matrix-free fashion.

Algorithm 4 Coupled Krylov method linear operator

1: function MULT(Y, X)	▷ Compute Jacobian-vector product with X
2: $(X_A, X_S) \leftarrow X$	▷ Extract aerodynamic and structural components
3: $Y_A \leftarrow \left(\frac{\partial A}{\partial w}\right)^T X_A$	▷ Evaluate diagonal contributions in parallel
$Y_S \leftarrow K^T X_S$	
4: $Y_A \leftarrow Y_A + \left(\frac{\partial F_A}{\partial w}\right)^T T Y_S$	▷ Add aerodynamic off-diagonal term
5: $Y_S \leftarrow Y_S + T^T \left(\frac{\partial X_V}{\partial X_S}\right)^T \left(\frac{\partial A}{\partial X_V}\right)^T X_A$	▷ Add structural off-diagonal term
6: $Y_S \leftarrow Y_S + T^T \left(\frac{\partial F_A}{\partial X_S}\right)^T T X_S$	▷ Add force $\frac{\partial S}{\partial u}$ contribution
7: end function	

4. Memory Requirements

The additional memory requirements for the fully coupled nonlinear and adjoint solution strategies are now considered. One advantage of the segregated Gauss–Seidel solutions strategies is that the coupled solution of the does not require any significant additional memory. This allows the multidisciplinary analysis and optimization to inherit the same memory footprint as the standalone solvers. If a Krylov method with a relatively long recurrence relationship is used to solve Equation (7) or Equation (14) significant additional memory is required. In the current framework, we use restarted FGMRES(m) from PETSc [13] to solve Equations (7) and (14) which requires $2m$ vectors of size \mathcal{R} to be stored in memory. For matrix-free disciplinary methods, these vectors require additional memory costs.

If Newton–Krylov methods are used to solve the disciplinary systems, however, it may be possible to reduce the memory requirement of the individual solvers when a coupled method is used. In the case of our CFD solver, when used stand-alone, a preconditioned GMRES(s) algorithm is used to approximately solve the Newton update equation and tightly solve the adjoint equation. The system is preconditioned using restrictive additive Schwartz and ILU(p) on each of the sub-domains, and requires s subspace vectors. For the coupled solution, it is possible to “exchange” some of the Krylov vectors used for discipline preconditioning to the outer Newton–Krylov or Krylov solver.

A fixed number of subspace vectors, n , can be used for the either the outer FGMRES solver of the discipline preconditioner. If we choose $m = \lfloor (n - s)/2 \rfloor$ the coupled methods and the Gauss–Seidel methods will use the same memory. In effect, we are transferring computational work from the discipline solvers to the coupled solvers. The effect is more outer iterations with less expensive inner iterations. The increase in the number of outer nonlinear or linear iterations also increases the computational cost of evaluating the off-diagonal terms. Eventually, these increasing cost leads to diminishing returns and there is a range of s values that lead to the best performance. This effect is investigated in Section IV.3.

IV. Results

1. Common Research Model Test Case

The configuration used for verification and optimization is the Common Research Model (CRM) [33] wing-body-tail configuration. It contains design features typical of a modern, wide-body, long range aircraft. This configuration was

analyzed extensively for the 4th AIAA Drag Prediction Workshop [34]. For our research, we created an structural model that conforms to the outer mould line provided that is representative of a modern airliner wingbox. The overall dimensions of the CRM is very similar to the Boeing 777-200ER and the publicly available Boeing 777 Airport Planning Document [35] provided the additional required data for the optimization problem such as the operational empty weight (OEW), maximum take-off weight (MTOW), the design range and design payload. Table 1 lists the key parameters used for the optimizations.

Table 1: CRM specifications

Parameter	Value	Units
Cruise Mach number	0.85	-
Cruise lift coefficient	0.5	-
Initial cruise altitude	35 000	ft
Span	58.6	m
Aspect ratio	9.0	-
Reference wing area	383.7	m ²
Sweep (Leading Edge)	37.4	°
Maximum take-off weight (MTOW)	298 000	kg
Operational empty weight	138 100	kg
Design range	7 725	nm
Design payload	34 000	kg
Reserve fuel	15 000	kg
Initial wing weight	29 200	kgf
Secondary wing weight	8000	kgf
Fixed weight	100 900	kgf
Thrust specific fuel consumption (c)	0.53	lb/(lbf · h)

2. Gradient Verification

Accurate sensitivity information is critical for the performance of any gradient-based optimization. All partial derivative terms for the coupled adjoint equation (14) and total sensitivity equation (15), are either computed with AD (using Taped [11]) or analytically. These terms can be evaluated at near machine precision allowing for extremely accurate sensitivity information. To ensure an entirely consistent formulation and implementation, we verify the coupled adjoint using the complex step method [36] as a benchmark applied to the entire system. We modified and created complexified versions of all the analysis modules that allow for the complex perturbation of any design variable and the complex perturbation and the resulting function of interest is the corresponding derivative, following the procedure outlined by Martins *et al.* [36]. The verification is carried out with the level 1 discretization given in Table 5. The multidisciplinary system and adjoint systems are converged to a relative tolerance of $\epsilon_{AS} = \epsilon_{SA} = 10^{-8}$. It is generally not possible to obtain any further convergence for the multidisciplinary system as the condition number of the structural Jacobian is typically $\mathcal{O}(10^9)$ for the flexible shell structures in our work. The resulting sensitivities are generally accurate to $\mathcal{O}(10^{-5})$. Obtaining more accurate results becomes difficult because very small discrepancies between the real and complex codes often appear, probably due to differences in compiler implementations and optimizations of certain built-in functions that do not always produce bit-compatible results. Given the highly ill-conditioned structural systems, these discrepancies are significant enough to prevent more accurate results. However, the level of accuracy achieved with the current coupled adjoint implementation is more than sufficient for gradient based optimization.

For the verification, we compute the derivative of two aerodynamic functions, C_L and C_D , and a structural function, the Kreisselmeier–Steinhauser (KS) stress aggregation function. The KS function aggregates the vonMises stress on the lower skin and lower stringers. Sensitivities of these three functions are computed with respect to global and local design variables. Global design variables explicitly affect the residual of all governing equations, while local variables only directly effect one discipline. The wing span and sweep, for example, are global variables, while α is a local aerodynamic variable, and the skin thickness of one lower surface skin panel at the root is local structural variable. Table 2 shows the resulting sensitivities as computed with the coupled adjoint method, the complex-step method and the finite-difference method. Complex-steps of $1 \times 10^{-40}j$ are used for the complex-step method and optimal step sizes for each variable are used for the finite-difference estimates.

The sensitivities computed with the coupled adjoint method differ from the complex-step method by relative errors of 10^{-3} to 10^{-6} . The functional sensitivities with respect to the local aerodynamic variable, α are typically more accurate due primarily to a more straightforward total derivative computation. The finite-difference sensitivities are

Table 2: Coupled adjoint sensitivity verification

Objective	Variable	Complex-step	Coupled adjoint		Finite-difference		
		Value	Value	Rel. Error	Value	Rel. Error	h_{opt}
C_L	Span	0.027265249	0.027266123	-3.21×10^{-5}	0.027858610	-2.18×10^{-2}	1×10^{-6}
	Sweep	-0.023601035	-0.023595480	2.35×10^{-4}	-0.023634861	-1.43×10^{-3}	1×10^{-4}
	α	0.143953101	0.143953092	5.96×10^{-8}	0.143632330	2.23×10^{-3}	1×10^{-6}
	t_{skin}	0.145264733	0.145266846	-1.45×10^{-5}	0.144828960	3.00×10^{-3}	1×10^{-5}
C_D	Span	0.000954129	0.000953880	2.61×10^{-4}	0.000911573	4.46×10^{-2}	1×10^{-5}
	Sweep	-0.001386782	-0.001388121	-9.65×10^{-4}	-0.001292728	6.78×10^{-2}	1×10^{-3}
	α	0.006336909	0.006336908	1.55×10^{-7}	0.006234932	1.61×10^{-2}	1×10^{-1}
	t_{skin}	0.007273797	0.007273570	3.12×10^{-5}	0.007141012	1.83×10^{-2}	1×10^{-6}
KS	Span	0.020412974	0.020414310	-6.55×10^{-5}	0.020939720	-2.58×10^{-2}	1×10^{-3}
	Sweep	-0.006399502	-0.006391973	1.18×10^{-3}	-0.005821129	9.04×10^{-2}	1×10^{-3}
	α	0.086843094	0.086843079	1.70×10^{-7}	0.092562735	-6.59×10^{-2}	1×10^{-3}
	t_{skin}	-0.660819304	-0.660817730	2.38×10^{-6}	-0.291121540	5.59×10^{-1}	1×10^{-5}

much less reliable and typically differ from the complex-step sensitivities by 10^{-1} to 10^{-3} . The limited accuracy of the finite-difference gradients limits their effectiveness in verifying a coupled adjoint implementation, as small errors of order $\mathcal{O}(10^{-2})$ cannot be easily detected.

3. Solution Methods Comparison

The two solution strategies for the solution of the aerostructural system and the two solution strategies for the adjoint system are compared. The comparison is carried out using the level 2 discretization as given in Table 5. The performance of the various methods at these mesh sizes is of particular interest, as we wish to perform the aerostructural optimization with this level of discretization. We consider two separate load cases: a 1 g cruise condition with only moderate elastic deformation and a 2.5 g maneuver condition with significantly more deflection and stronger aerostructural coupling. Several variations of the nonlinear block Gauss–Seidel (NLBGS) and Newton–Krylov (NK) methods are considered. A total of four fixed under-relaxation are used, namely 0.25, 0.5, 0.75 and 1.0 (no under-relaxation). For these analysis $\theta^{(k)} = \theta^{(0)}$ for all iterations. An additional analysis with Aitken acceleration ($\theta^{(0)} = 0.5$) is also computed. For all NLBGS analysis, $\epsilon_A = 0.1$ and $\epsilon_S = 1 \times 10^{-3}$. However, since we are using a direct factorization method for the structural analysis, a lower error than ϵ_S is normally achieved.

For the NLBGS analysis, the aerodynamic Newton–Krylov solver uses GMRES(75). This system is preconditioned with additive Schwartz [37] (overlap 1) and ILU(2) on each of the sub-domains. For the Newton–Krylov algorithm, three non-Aitken accelerated relaxed Gauss–Seidel start-up iterations are performed for the 1 g case and five for the 2.5 g case. Three variants of the algorithm are evaluated, each with a different subspace size for the aerodynamic part of the coupled block-Jacobi preconditioner. To ensure a fair comparison using constant memory, the size of the FGMRES outer subspace is taken as $m = \lfloor (75 - s)/2 \rfloor$, which yields similar memory usage for the CFD processors. While the structural processors do have additional memory requirements, this typically isn't a limiting factor for our problems and is considered acceptable.

Table 3 lists the timing results for each of the test combinations. The N_{iter} column is the number of iterations required for convergence, which has a slightly different meaning, depending on the solution method. For the NLBGS methods, this is the number of iterations required for convergence. For the NK solution methods, it refers to the total number of nonlinear residual evaluations. Since the cost of the preconditioned nonlinear residual evaluations varies with the strength of the preconditioned (value of s) the number of iterations is not directly comparable across methods. A “—” symbol represents parameters that are not relevant.

A number of overall trends can be determined from the results. For both load cases, fixed relaxation parameters of 0.25 or 1.0 consistently performed poorly, while the values of 0.5 and 0.75 yielded reasonable performance. The lowest under-relaxation of 0.25 is very robust, but takes a large number of iterations to converge. Conversely, without under-relaxation, the $\theta = 1.0$ case consistently over and under shoots the actual structural displacement, especially for the 2.5 g condition, leading to a large number of iterations. The Aitken acceleration scheme aids the Gauss–Seidel scheme considerably, yielding the best NLBGS solution time for the 2.5 g load case and the second best GS solution

Table 3: Nonlinear solution method comparison

Solver	$\theta^{(0)}$	Aitken accel.	Inner size, (s)	Outer size, (m)	N_{iter}	CFD time, (s)	Warp time, (s)	CSM time, (s)	Total time, (s)
1.0 g Load factor									
NLBGS	0.25	N	—	—	49	384.4	6.2	23.0	395.6
NLBGS	0.50	N	—	—	22	218.7	2.8	20.3	223.7
NLBGS	0.75	N	—	—	16	190.3	2.0	19.7	193.9
NLBGS	1.00	N	—	—	22	216.8	2.8	21.8	221.9
NLBGS	0.50	Y	—	—	18	200.3	2.3	19.9	204.5
NK	0.50	N	5	35	202	207.2	25.3	32.5	239.5
NK	0.50	N	10	32	79	162.4	9.8	24.1	175.5
NK	0.50	N	20	27	56	183.6	3.9	22.2	193.0
2.5 g Load factor									
NLBGS	0.25	N	—	—	52	349.9	6.6	23.2	358.9
NLBGS	0.50	N	—	—	28	254.9	3.5	20.9	261.3
NLBGS	0.75	N	—	—	25	272.2	3.1	20.6	278.0
NLBGS	1.00	N	—	—	52	363.0	6.6	23.2	375.0
NLBGS	0.50	Y	—	—	25	263.1	3.1	20.6	268.8
NK	0.50	N	5	35	225	234.8	28.2	34.8	270.8
NK	0.50	N	10	32	118	214.6	14.7	26.8	233.8
NK	0.50	N	20	27	55	193.9	6.8	22.3	203.2

time for the 1 g case. Given the ease of implementation and increase in performance, the Aitken-accelerated form of the Gauss–Seidel method, is the most efficient of the segregated analysis methods.

The Newton–Krylov method generally performs better than the Gauss–Seidel methods and the lowest solution time for each of the two load cases comes from NK methods. The effect of the strength of the inner preconditioner on the number of inner iterations is evident in Table 3. As expected, $s = 5$ requires the most, outer iterations, with each iteration having the lowest computational cost. With a very large number of outer iterations, the demand on mesh movement scheme increases considerably, requiring over 200 mesh perturbations. Since our scheme is very efficient — requiring approximately 0.1 seconds per movement — a small inner subspace is still viable. However, better solution times are obtained using a slighter stronger preconditioner. Solution times with $s = 10$ and $s = 20$ are similar and suggest that for computational problems with this discretization level, which is the case in our current framework, 10–20 is a reasonable range for this parameter. For the 1 g load case, the best NK solution time is 9% lower than the best NLBGS time, and 22% faster for the 2.5 g load case. Not surprisingly, the advantage of the coupled solution technique is more pronounced in the more strongly coupled 2.5 g load condition.

We also perform a similar comparison for the coupled adjoint system. For the LBGS solver, two fixed damping ratios, 0.5 and 1.0 are considered, as well as three relative forcing tolerances for the aerodynamic adjoint solution: $\epsilon_A=0.1, 0.25$ and 0.5 . The same three variants of the coupled solver are considered with inner aerodynamic subspace sizes of 5, 10, and 20. As with the nonlinear solution comparison, the memory requirements are held fixed: The LBGS solver uses GMRES(75) while the coupled solver uses GMRES(s) for the aerodynamic preconditioner and FGMRES($\lfloor(75 - s)/2\rfloor$) for the outer subspace. The results are given in Table 4.

The benefits of the monolithic solution method are even more compelling for the adjoint solution. As is the case with the nonlinear solution, all Gauss–Seidel solutions were able to converge, albeit with some difficulty. For 1 g case the a damping factor of 1.0 yielded faster solution times, while for the 2.5 g case, the damping factor of 0.5 was generally faster. Both test cases converged faster with relatively weak forcing tolerances of only 0.5. Additional accuracy for each approximate solution is unnecessary and only increases computational cost. The wide variation in the solution times for the LBGS method indicates a strong dependence on both the forcing tolerances and the damping ratio. All three variants of the coupled Krylov solver converged without issue and are typically significantly faster than best LBGS combination. The same observation regarding the small inner subspace sizes ($s \approx 5$) made for the nonlinear solution applies to the adjoint solution but to an even greater extent. The warping derivatives, while efficient, are more costly than their forward equivalent. The cumulative time required for the warping derivatives represents a much larger proportion of the total solution time for all of the coupled solution results. For the $s = 5$ case, it represents

Table 4: Adjoint solution method comparison

Solver	θ	Forcing tolerance	Inner size, (s)	Outer size, (m)	N_{iter}	CFD time, (s)	Warp time, (s)	CSM time, (s)	Total time, (s)
1.0 g Load factor									
LBGS	0.5	0.10	—	—	21	542.6	9.1	2.1	554.0
LBGS	0.5	0.25	—	—	22	304.2	9.5	2.3	316.1
LBGS	0.5	0.50	—	—	23	141.7	10.0	2.3	154.1
LBGS	1.0	0.10	—	—	17	375.9	7.4	1.7	385.0
LBGS	1.0	0.25	—	—	16	158.9	7.0	1.6	167.6
LBGS	1.0	0.50	—	—	21	111.6	9.1	2.1	123.0
CKM	—	—	5	35	120	78.2	53.4	10.6	136.3
CKM	—	—	10	32	58	71.3	25.1	5.1	99.6
CKM	—	—	20	27	40	99.8	17.3	3.5	118.5
2.5 g Load factor									
LBGS	0.5	0.1	—	—	21	552.4	9.1	2.1	563.8
LBGS	0.5	0.25	—	—	21	311.1	9.2	2.1	322.5
LBGS	0.5	0.50	—	—	24	139.2	10.5	2.4	152.2
LBGS	1.0	0.10	—	—	79	4518.4	34.5	7.8	4561.3
LBGS	1.0	0.25	—	—	42	308.1	18.2	4.2	330.1
LBGS	1.0	0.50	—	—	46	175.9	20.0	4.5	200.8
CKM	—	—	5	35	125	81.2	55.5	11.0	141.8
CKM	—	—	10	32	53	77.6	27.4	5.6	107.4
CKM	—	—	20	27	47	115.9	20.4	4.1	137.9

approximately 40% of the overall time and approximately 15% for the $s = 20$ case. Values of $s = 10$ yielded the best solution times. However we caution the optimal values of s are problem dependent and may also depend on other solution parameters such as the level of fill used for the aerodynamic preconditioner. Generally, s values ranging from 10–20 give fast, robust convergence, similar to the performance of the nonlinear solution method.

Comparing the best LBGS solution time with the best CKM times shows a 19% reduction in time for the 1 g load case and 29% reduction for the 2.5 g case. We conclude for the type of static aerostructural deformations we expect to find on a subsonic transport aircraft wing, the fully coupled solution methods (nonlinear and adjoint) can yield computational savings in the order of 10–30%.

4. Parallel Scalability Study

We now examine the overall parallel scalability of the coupled solution and adjoint techniques. We consider a sequence of three discretization levels, where each level doubles the spatial resolution of the previous level. For the CFD solver each level increases the number of cells by a factor of eight. Since the structural solver uses shell elements, doubling the spatial resolution results in four times the number of elements and approximately four times the number of structural degrees of freedom. Table 5 lists the sizes of each discretization.

This weak scaling study attempts to keep the number of degrees of freedom per processor roughly constant. For the CSM problem, near perfect load balancing is achieved for each mesh level, but this is not the case for the CFD problem. The load imbalance is computed according to: $\text{Imbalance} = \max(N_i) / N_{\text{even}} - 1$, where N_i is the number of cells on each processor and $N_{\text{even}} = N_{\text{total}} / N_P$ is the number of cells resulting from a perfect balance balancing. The load imbalance of for mesh level 3 indicates at least 1 critical processor is assigned 25% more computational work than is required on average, leading to computation times 25% longer than would result from perfect load balancing.

Solutions are generated for a Mach number of 0.85, a fixed angle of attack of 2.0° and static pressure of 30 000 Pa resulting in approximately a 1 g load condition. The nonlinear equations are solved using the NLBGS method, while the adjoint equations are solved with the coupled Krylov method. The adjoint solution also evaluates the sensitivity of lift with respect to 476 design variables. Both the nonlinear and adjoint solutions are converged to relative tolerance $\epsilon_{AS} = \epsilon_{SA} = 10^{-6}$. Figure 3 shows the C_p contours for each converged solution as well as the fraction of yield stress on the upper surface of the wing box. A section of the computational meshes in each discipline is also provided for

Table 5: Mesh sizes

Level	CFD				CSM			
	Cells	DOF	N_P	Imbalance	Elements	Nodes	DOF	N_P
1	263 762	1 318 810	8	1.5%	13 512	13 522	81 132	2
2	2 110 096	10 550 480	64	13.1%	54 048	54 030	324 180	8
3	16 288 0768	84 403 840	512	25.2%	216 192	216 118	1 296 708	32

visual reference. Comparisons of key aerodynamic and structural performance results are given in Table 6.




Figure 3: CFD and CSM solution comparison for each mesh level

The general characteristics of the solutions at each mesh level are similar, but a closer investigation reveals several differences. As the CFD mesh is refined, the normal shock on the upper surface of the wing becomes more pronounced, leading to a more accurate prediction of the wave drag. The lift coefficient is predicted reasonably well on all three mesh levels, while the drag coefficient reduces considerably as the mesh is refined due to the reduction of spurious drag caused by the artificial dissipation scheme. As expected, the higher resolution structural mesh has larger tip displacement and higher stress levels.

We now examine the scalability of our methods by examining the time required to compute each solution and

adjoint. These computations as well as all other results in this paper were performed using Intel Xeon E5540 processors connected with a 4x-DDR non-blocking InfiniBand fabric interconnect [38]. A detailed timing breakdown of the sensitivity analysis performed at the above converged solution is given in Table 7. The fraction column is nondimensionalized relative to the aerostructural analysis time for that level.

The aerostructural solution times for all three mesh levels show that the majority of the solution time is taken by the CFD solution. The mesh movement scheme comprises only 2–3% of the total solution time. Similarly, the time required for transferring the loads and displacement is very low. We would expect the total solution time for a NLBGS solution to the sum of the individual disciplinary solution times. However, since we are using a direct solution method for the CSM, the majority of solution time is constituted by the matrix factorization, which we can overlap with the first CFD Gauss–Seidel iteration.

For the coupled adjoint solution, the one time assembly costs account for the time required to compute the aerodynamic state residual matrix, $\partial\mathcal{A}/\partial w$, the spatial residual matrix, $\partial\mathcal{A}/\partial X_V$, and the two coupling matrices $\partial F_A/\partial w$ and $\partial F_A/\partial X_S$. This amounts to approximately 35% of the analysis time for the level 2 mesh, but is quickly amortized over multiple adjoint solutions. Since the stiffness matrix for the structural discipline is symmetric, the original factorization for the solution method is reused. The bulk of the coupled adjoint solution time is spent in each of the disciplinary solvers, which amounts to approximately 35% of the nonlinear solution time. The computational cost for the off-diagonal terms is dominated by the time required to compute $\partial\mathcal{A}/\partial u$, listed in Table 7 as “Mesh sensitivities”. This typically amounts to the 20–30% of the coupled adjoint solution time for the level 1 and 2 mesh. The time required for the mesh sensitivities on the larger mesh is disproportionately large. The reason for this however, is due to load balancing of the mesh movement algorithm. The mesh movement algorithm currently does not support block splitting, which means once the number of processors exceeds the number of blocks in the mesh, (144 in this case) the remaining processors are idle during mesh operations. If the level 3 mesh contained a correspondingly larger number of blocks, this behaviour would not be observed. The total sensitivity time includes the calculation of all partial derivative terms in the total sensitivity equation.

It is instructive to examine how the convergence characteristics of the nonlinear solution and linear adjoint solutions change as the size of the computation problem increases by nearly two orders of magnitude. Figure 4 shows the nonlinear convergence and Figure 5 gives the convergence for the adjoint.

For both the level 1 and 2 meshes convergence to 1×10^{-6} is achieved in approximately the same number of iterations (16 and 19 respectively), while engineering accuracy (1×10^{-3}) for the lift to drag ratio is achieved in approximately 10 iterations. However, the level 3 solution requires significantly more (36) and engineering accuracy is not achieved until iteration 20.

Moving from mesh level 2 to 3, the number of NLBGS iterations have doubled but the wall time has increased by a factor of 4.3. Since a fixed aerodynamic forcing tolerance is used, each iteration is also more costly on the larger mesh. The performance for the coupled adjoint solution is similar. In this case, the cost of each iteration is similar for all three mesh levels such that the number of iterations required for convergence in Figure 5 is representative of the overall solution time. All of the adjoint solutions use the same aerodynamic preconditioning settings, ILU(1) and additive Schwartz(1), resulting in nearly constant memory usage across the mesh levels. For mesh level 3, faster convergence times can be achieved by using stronger preconditioning which reduces the condition number of the preconditioned system.

Table 6: Aerostructural scaling results

Quantity	Level 1	Level 2	Level 3
C_L	0.4871	0.4960	0.4968
C_D	0.0186	0.0121	0.0109
C_M	-0.0250	-0.0224	-0.0216
Tip displacement, m	1.54	1.61	1.62
Lower skin/stringer KS	0.328	0.365	0.491
Upper skin/stringer KS	0.459	0.492	0.723
Rib/spar KS	0.367	0.431	0.512
Nonlinear Gauss–Seidel iterations	16	19	36
Coupled Krylov adjoint iterations	24	46	117

Table 7: Coupled adjoint timings

	Level 1		Level 2		Level 3	
	Time, s	Fraction	Time, (s)	Fraction	Time, (s)	Fraction
Aerostructural solution	117.1	1.000	209.7	1.000	897.2	1.000
CFD solution	112.6	0.962	205.1	0.978	863.2	0.962
CSM solution	7.24	0.062	20.0	0.095	51.1	0.057
Mesh movement	1.8	0.015	2.4	0.011	24.3	0.027
Coupled adjoint solution	53.2	0.454	112.7	0.537	579.6	0.646
Setup	42.9	0.366	71.3	0.340	52.6	0.059
CFD solution	43.7	0.371	90.1	0.430	241.3	0.269
CSM solution	0.9	0.008	4.0	0.019	19.3	0.022
Mesh sensitivities	8.9	0.076	20.7	0.097	328.4	0.366
Total sensitivity	5.4	0.046	5.7	0.027	7.0	0.012
First gradient		0.823		0.913		0.716
Subsequent gradients		0.457		0.573		0.657

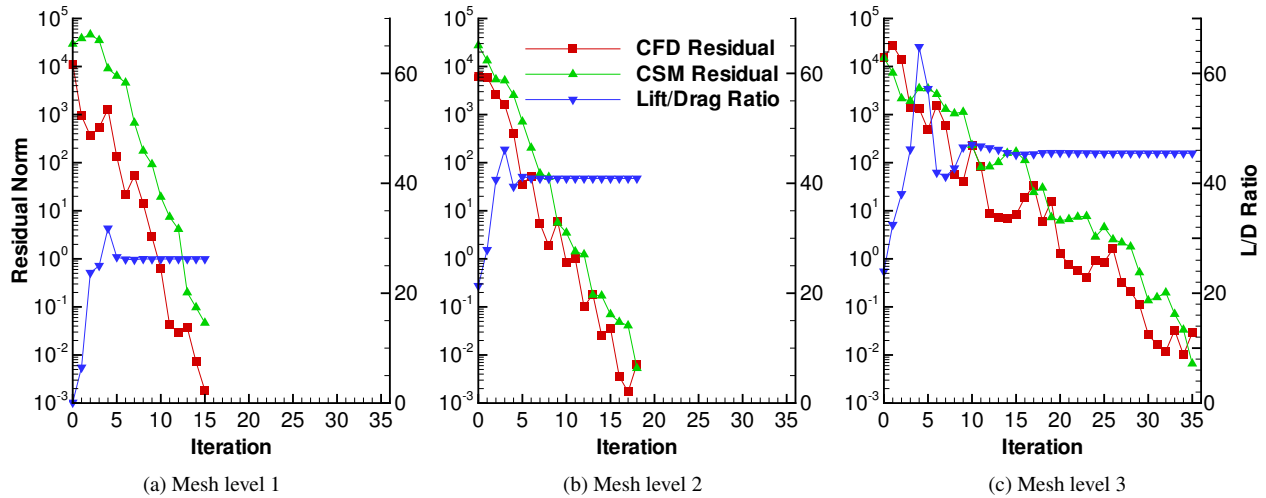


Figure 4: NLBGS convergence for the three mesh levels

5. Optimization Results

This section presents the results from two optimizations performed using the methods described in Sections II and III. Two multidisciplinary objective functions are considered: A minimization of the Maximum Takeoff Weight (MTOW) and a minimization of the fuel burn. The goal is to make the optimization problem formulation as realistic as possible.

5.1. Design and Maneuver Conditions

The operating conditions used for the optimizations are now described. In order to evaluate the MTOW or fuel burn objectives, we estimate the weight at the start of cruise using the Breguet range equation. The only aerodynamic input to this computation is the lift to drag ratio of the aircraft. While a single analysis point is sufficient to estimate the lift to drag ratio, single points optimizations, (especially in the transonic regime) can produce very optimal performance at a single operating point at the cost of significant degradation in other important off-design conditions [39]. To help alleviate this problem, we consider a multi-point aerostructural optimization that considers the average performance over multiple flight conditions which will hopefully lead to a more robust design. We consider a total of five cruise operating conditions centered about the design operating design condition for the CRM configuration, which is $M=0.85$, $C_L=0.5$. The first flight condition, matches the CRM design condition closely at the chosen flight altitude. The second




Figure 5: Convergence history for lift adjoint using the coupled Krylov adjoint solver

and third design points use the same altitude but slightly lower and higher mach numbers respectively. The final two flight conditions, use the design Mach number of 0.85, but at ± 1000 ft of altitude, which has the effect of decreasing and increasing the flight C_L .

Two separate maneuver conditions are considered: a 2.5 g symmetric pull-up maneuver and a 1.3 g acceleration due to gust. Maneuver condition 1 is designed to represent a 2.5 g limit load for the wing structure. The second maneuver condition is designed to simulate a gust load on the aircraft in a cruise condition. The purpose of this maneuver condition is explained in the discussion of the optimization problem constraints.

A final analysis point is also required for the estimate of the aircraft's static margin. Using only the steady state tools we have developed, it is possible to estimate the static margin of the deformed configuration using the following formula

$$K_n = -\frac{C_{M_\alpha}}{C_{L_\alpha}}. \quad (23)$$

Further details on approximating static and dynamic stability derivatives for high fidelity CFD optimization can be found in Mader *et al.* [40]. The derivatives C_{M_α} and C_{L_α} are estimated using a forward finite-difference with step size $\Delta\alpha = 0.1^\circ$. The analysis for cruise condition 1 serves as the base value and the additional stability point provides the additional analysis to complete the derivative calculations.

Table 8: Operating conditions

Group	Identifier	Mach	Altitude, (ft)	Load Factor
Cruise	1	0.85	35 000	1.0
	2	0.84	35 000	1.0
	3	0.86	35 000	1.0
	4	0.85	34 000	1.0
	5	0.85	36 000	1.0
Maneuver	1	0.86	20 000	2.5
	2	0.85	32 000	1.3
Stability	1	0.85	35 000	1.0

5.2. Objective Function

The choice of objective function has a significant impact on the resulting optimized design. The objective function must be carefully chosen to accurately reflect the design intent for a particular aircraft. In many ways, this is a much more challenging problem for aerostructural analysis than it is for either aerodynamics or structures separately. As an

example, consider a typical multidisciplinary objective, maximizing the range using Breguet range equation:

$$R = \frac{V}{c} \frac{C_L}{C_D} \ln \left(\frac{W_1}{W_2} \right). \quad (24)$$

Here, R is the range, V is the flight speed, c is the thrust-specific fuel consumption and W_1 and W_2 are the initial and final cruise weights, respectively. We assume that the aircraft is able to climb continuously as fuel is consumed maintaining the same lift to drag ratio and Mach number. If the Range equation is used as an objective for aerodynamic optimization, the ratio of the cruise weights remains fixed, and assuming a constant Mach number and thrust specific fuel consumption, the range equation reduces to an L/D maximization or lift constrained drag minimization. This simplification removes the need to know the design range of the aircraft. Similarly, if we consider a structural optimization, only the weight ratio remains which leads to a common mass minimization problem.

In both cases the single discipline optimization problems lead to simplified formulations and, critically, ignores the *magnitude* of the terms that are assumed constant. An aerostructural analysis can manipulate the L/D and $\ln(W_1/W_2)$ terms simultaneously, and it is critical that the magnitude of these terms reflect the intent of the designer.

A quantitative analysis can be performed to determine the tradeoff between drag and weight reduction and answers the the question: How many kilograms of structural weight is equivalent to one drag count? Consider the partial derivatives of the range objective with respect to the drag coefficient C_D (γ) and the structural weight W_2 (β) as well as their ratio:

$$\gamma = \frac{\partial I}{\partial C_D} = -\frac{V}{c_T} \frac{C_L}{C_D^2} \ln \left(\frac{W_1}{W_2} \right) \quad (25)$$

$$\beta = \frac{\partial I}{\partial W_2} = -\frac{V}{c_T} \frac{C_L}{C_D} \frac{1}{W_2} \quad (26)$$

$$\frac{\gamma}{\beta} = \frac{\partial W_2}{\partial C_D} = \frac{W_2}{C_D} \ln \left(\frac{W_1}{W_2} \right). \quad (27)$$

The last expression shows us the weight-drag tradeoff scales with the final cruise weight, the natural logarithm of the weight ratio, and is inversely proportional to the cruise drag coefficient. Unlike the single discipline optimization cases where a simple reduction of drag or weight was sufficient, the multidisciplinary objective requires knowledge of the actual aerodynamic performance (C_D) and actual structural performance ($\ln(W_1/W_w)$) to determine the correct multidisciplinary trades. This places additional burden on the analysis as we must include *all* components of the drag and weight even if they are not explicitly modeled to achieve the correct aerostructural trades.

For our two objectives, we use the Breguet range Equation (24), fix the desired range, R , and compute the weight at the start of cruise W_1 . The fuel burn is then $W_1 - W_2$. For the purposes of the optimization problems, we only consider the fuel burn during a single cruise segment and ignore the fuel burned during take-off, climb and descent. For this reason we assume $W_1 = \text{MTOW}$. Using a similar linearization analysis that we performed for the Range objective we find:

$$I = W_1 = \text{MTOW} \quad (28)$$

$$\gamma = \frac{\partial I}{\partial C_D} = e^{\frac{C_T R C_D}{V C_L}} W_2 \left(\frac{C_T R}{V C_L} \right) = \frac{W_1}{C_D} \ln \left(\frac{W_1}{W_2} \right) \quad (29)$$

$$\beta = \frac{\partial I}{\partial W_2} = e^{\frac{C_T R C_D}{V C_L}} = \frac{W_1}{W_2} \quad (30)$$

$$\frac{\gamma}{\beta} = \frac{\partial W_2}{\partial C_D} = \frac{W_2}{C_D} \ln \left(\frac{W_1}{W_2} \right) \quad (31)$$

$$I = W_1 - W_2 = \text{Fuel Burn} \quad (32)$$

$$\gamma = \frac{\partial I}{\partial C_D} = e^{\frac{C_T R C_D}{V C_L}} W_2 \left(\frac{C_T R}{V C_L} \right) = \frac{W_1}{C_D} \ln \left(\frac{W_1}{W_2} \right) \quad (33)$$

$$\beta = \frac{\partial I}{\partial W_2} = e^{\frac{C_T R C_D}{V C_L}} - 1 = \frac{W_1}{W_2} - 1 \quad (34)$$

$$\frac{\gamma}{\beta} = \frac{\partial W_2}{\partial C_D} = \frac{W_1 W_2}{C_D (W_1 - W_2)} \ln \left(\frac{W_1}{W_2} \right). \quad (35)$$

First, we find the MTOW objective and Range objectives yield identical linearizations, with the fuel burn objective distinct from the other two. Additionally, the fuel burn objective sensitivity differs from the MTOW sensitivity by a

factor of $W_1 / (W_1 - W_2) = W_1 / (\text{Fuel Burn})$ which is just the reciprocal of the fuel fraction. Since the fuel fraction is always less than one, the γ/β ratio will always be higher than the ratio for MTOW. A higher γ/β places a higher weight equivalent to a single drag count. Thus, a higher γ/β value favor more aerodynamically-biased solutions trading increased weight for aerodynamic efficiency, while lower γ/β values favors more structurally-biased solutions with reduced structural weight over aerodynamic efficiency. Using the data in Table 1 we can produce estimates for the objective sensitivities as well as the γ/β ratio.

Table 9: Objective sensitivities

		MTOW	Fuel Burn
γ	$= \frac{\partial I}{\partial C_{D_{\text{count}}}}$	497.0 kg	497.0 kg
β	$= \frac{\partial I}{\partial W_2}$	1.59	0.59
γ/β	$= \frac{\partial W_2}{\partial C_{D_{\text{count}}}}$	312.6 kg	790.2 kg

The fuel fraction for the CRM aircraft at design payload and range is approximately 0.37, which indicates the ratio of γ/β for the two objective is $1/0.37 = 2.69$. With the MTOW objective, 1 drag count is worth approximately 313 kg but for the fuel burn objective 1 drag count is equivalent to 790 kg. The effect of the multidisciplinary objective on the aerostructural optimization is discussed in Section 5.7.

To compute the composite objective function for the multipoint analysis, we analyze each design point at the weight corresponding to the midpoint of the mission. The final cruise weight, W_2 , is computed according to

$$W_2 = M + \text{Fixed Weight} + \text{Reserve Fuel Weight} + A/A_{\text{ref}} \times \text{Secondary Wing Weight} \quad (36)$$

where M is the weight of the primary wing structure, A is the projected wing area and A_{ref} is the initial projected wing area. The last term is used to account for the increasing non-structural weight of the wing as the area increases. For each condition, the L/D ratio along with the Breguet range equation is used to determine the weight at the start of cruise, W_1 . For the optimizations considered, we apply a simple average over each cruise condition to formulate the composite objectives.

5.3. Design Variables

The two optimizations presented use hundreds of design variables to parametrize the aerodynamic and structural models. As is typical in a multidisciplinary analysis, we can split the design variables into global variables, which directly affect more than one discipline and local variables, which only affect a single discipline. Table 10 lists all the optimization variables.

Table 10: Optimization design variables

Global design variables		Aerodynamic design variables		Structural design variables	
Description	Quantity	Description	Quantity	Description	Quantity
Span	1	Angle of attack	7	Upper skin	54
Sweep	2	Tail rotation angle	7	Lower skin	54
Chord	4	Mean aerodynamic chord	1	Upper stringers	54
Twist	5	X_{CG}	1	Lower stringers	54
Shape	160			Ribs and stiffeners	36
				Spars	36
Total	172	Total	16	Total	288
				Grand total	476

We use a free from deformation (FFD) volume approach to make geometric perturbations to the geometry. Further information on our approach can be found in Kenway *et al.* [17]. The main wing planform variables, span, sweep, chord and twist are all global variables, as they directly affect the geometry in each discipline. A graphical representation of the geometric design variables, as well as the level 2 mesh discretization used for the optimization is shown in Figure 6. Chords are modified at the root, yehudi break, near tip and tip sections with the remaining sections interpolated. Five twist angles are defined similarity. Two sweep variables are specified: the first, at the second last

station, sweeps the entire leading edge of the which, by construction is constrained to remain straight. The second tip variable only moves the outer tip section. The shape variables are used to perturb the coefficients of the FFD volume surrounding the wing in the z (normal) direction. These shape variables prescribe the chord and span-wise airfoil thicknesses distribution and no additional explicit thickness variables are required.




Figure 6: Geometric design variables and FFD volumes (left) and CFD surface mesh (right) for the CRM design optimization problem.

The structural skin variables are grouped in with 18 variables in the span-wise and 3 variables in the chord-wise direction respectively. The stringer variables are grouped similarly. The ribs and rib stiffeners each have 18 variables in the span-wise direction as do each of the leading and trailing edge spars. Each of the five cruise conditions and two maneuver conditions have independent angle of attach and tail rotation angles to provide the required degrees of freedom to meet both the lift and moment constraints. Finally, MAC and X_{CG} are target variables that simplifies our specific implementation. The initial values of geometric and aerodynamic design variables are chosen to regenerate exactly the original CRM geometry. Without a nominal baseline structure we must determine a reasonably efficient initial structural model to enable comparisons between the initial and optimized designs. The design of the initial structure is generated by performing a stress constrained mass minimization optimization. The fixed loads are derived from an aerostructural analysis and the same 6 KS stress constraints described in Section 5.4 are used to ensure a feasible solution. The resulting optimal design, is then used as the initial structural design for the two optimization problems. The result of this structural optimization can be seen in Figures 14c and 14d. Figure 7 shows the internal layout of the structure as well as the grouping of the structural design variables.

5.4. Design Constraints

In this section we describe the constraints that are used for the MTOW and fuel burn optimizations. For simplicity, we divide the constraints into three groups: geometric and target constraints, aerodynamic constraints and structural constraints as listed in Table 11. The first two geometric constraints t_{LE} and t_{TE} are used to constrain the initial wing thickness at 2.5% and 97.5% chord. This effectively constrains the leading edge radius and prevents crossover near the sharp trailing edge. The purpose of constraining the leading edge radius is twofold: First, it eliminates excessively sharp leading edges, that are difficult to manufacture, and second, helps ensure the high-speed aerostructural optimiza-




Figure 7: Structural design variable groups

Table 11: Optimization constraints

Geometric/target constraints			Aerodynamic constraints			Structural constraints		
Description	Quantity		Description	Quantity		Description	Quantity	
t_{LE}/t_{LEinit}	\geq	1.0 11	Cruise: $L - W$	$=$	0.0 5	2.5 g Lower skin: KS	\leq	1.0 1
t_{TE}/t_{TEinit}	\geq	1.0 11	Cruise: C_{m_y}	$=$	0.0 5	2.5 g Upper skin: KS	\leq	1.0 1
A/A_{init}	\geq	1.0 1	Maneuver: $L - W$	$=$	0.0 2	2.5 g Rib/spars: KS	\leq	1.0 1
V/V_{init}	\geq	1.0 1	Maneuver: C_{m_y}	$=$	0.0 2	1.3 g Lower skin: KS	\leq	0.42 1
$t_{TE\ Spar}$	\geq	0.20 5	Static margin: K_n	\geq	0.15 1	1.3 g Upper skin: KS	\leq	1.0 1
$t_{Tip}/t_{Tipinit}$	\geq	0.5 5				1.3 g Rib/spars: KS	\leq	1.0 1
$MAC - MAC^*$	$=$	0.0 1						
$X_{CG} - X_{CG}^*$	$=$	0.0 1						
Total		36	Total		15	Total		6
						Grand total		57

tion does not significantly impact the low speed $C_{L_{max}}$ performance, which is largely governed by the roundness of the leading edge. We constrain the projected wing area to be no less than the initial area. This constraint is used to ensure adequate low speed performance, more specifically the takeoff distance and approach speed. Even with a fairly complete structural model, a volume constraint is required. The volume constraint is only computed over the spar-box and is used to ensure that the optimized wing is able to carry at least the same amount of fuel as the initial design.

Several additional thickness constraints are also enforced. Minimum trailing edge spar height constraints, $t_{TE\ spar}$, are used over the outboard section of the wing. This constraint is intended to ensure adequate vertical space is available for the trailing edge actuation devices for the flaps and ailerons. The final thickness constraint, t_{tip} , is used to ensure the optimization does not produce an unrealistically thin wing tip.

Each cruise and maneuver condition have a constraint on the total lift, L , as well as on the pitch moment coefficient C_M . The static margin of the first cruise condition is constrained to be greater than 15%. The reference point for the moment computation is taken to be 25% of the mean aerodynamic chord (MAC), which changes with the planform design variables. By including a pitch moment constraint as well as a constraint on the full configuration static margin, we allow the optimization to trade drag reduction from aft-loaded super-critical profiles with the induced drag penalty required to trim the configuration. These constraints ensure the optimal designs produce the required lift, with zero pitch moment and have sufficient static stability.

Lastly, we must constrain the computed stress on the structure. The level 2 structural model used for the optimization consists of over 50 000 elements. Constraining the stress in each of these elements would require the solution of the corresponding number of coupled adjoint vectors. To reduce the number of required adjoint solutions, we use

Kreisselmeier–Steinhauser (KS) constraint aggregation technique [32]. Each maneuver condition uses three KS functions: The first for the lower wing skin and stringers, the second for the upper wing skin and stringers and the third for the spars, ribs and rib stiffeners. The compression members in the upper wing skin and stringers are assumed to be Aluminum 7050 with a maximum allowable stress of 300 MPa. The remainder of the primary wing structure is assumed to be manufactured from Aluminum 2024 with a maximum allowable stress of 324 MPa.

For the first maneuver condition, the maximum vonMises stress must be below the limiting stress, which requires the three KS functions to be less than 1.0. The second maneuver constraint is used to emulate a fatigue criteria for the lower wing skin and stringers that are under tension during normal loading. We limit the stress on the lower wing skin and stringers to be below 138 MPa for the 1.3 g load condition. The upper limit for this KS function is 0.42. The remaining two KS functions for the 1.3 g maneuver condition retain a the maximum KS value of 1.0.

5.5. Optimization Algorithm

For design optimization problems with hundreds of design variables (476 in our case), a gradient based optimization algorithm is the only viable option. For this work we use SNOPT [41], an optimizer based on the SQP approach. SNOPT is well suited for large scaled, constrained nonlinear optimization problems. Since the coupled adjoint method allows us to compute gradients of a function of interest in less time than a nonlinear function, this gradient based method significantly reduces the computational effort required to find an optimal solution.

5.6. Computational Resources

The two optimizations are performed on a massively parallel supercomputer [38]. Each optimization function evaluation requires the solution of eight flow solutions: five for the cruise conditions, two for the maneuver conditions and one for the stability point. To reduce the wall time required for the optimizations, each of these parallel analysis are carried out concurrently. The cruise condition and stability condition use the level 2 discretization meshes from Table 5 with slightly differing numbers of processors: 52 for the CFD problem and 4 for the CSM problem. The two maneuver conditions use a smaller CFD mesh with 1.2 million cells with the same level 2 structural discretization. The maneuver conditions use 45 CFD processors and 4 CSM processors. In order to ensure good overall computational efficiency for the optimization problem, we must ensure the time required to run each concurrent analysis is approximately equal. This further complicated by fact that the cruise analysis points require the computation of three adjoint vectors (lift, drag and moment), the maneuver conditions require five adjoints, (lift, moment, and three KS functions) and the stability condition requires two (lift and moment). The selected processor counts achieve good discipline-level load balancing (especially for the CFD problem) as well as good optimization system-level load balancing. The total number of processors used the optimization is $435 = 56 \times 5 + 49 \times 2 + 56 + 1$. The single additional processor at the end is used for the skin friction computation. The MTOW optimized required 34.5 hours of wall time while the fuel burn optimization required 36.5 hours.

5.7. Optimization Results Discussion

We now examine the optimization results from two optimization: an MTOW minimization and a fuel burn minimization. The results for each optimization are presented concurrently as to highlight the effects of the chosen objective function.

First we examine the optimization convergence history for each problem which is given in Figure 8. Feasibility is the maximum constraint violation and is a measure of how closely the non linear constraints are satisfied. Optimality refers to the norm of the augmented Lagrangian function used in SNOPT. Additionally the evolution of each objective is also given. 150 major optimization iterations are performed for each optimization which results in approximately two orders of magnitude reduction in the optimality. A lower optimality indicates a more accurate solution to the Karush–Kuhn–Tucker (KKT) conditions and results in a more well defined local minima. We have found these complex aerosturctural optimization problems converge slowly and only produce weak local minima. However, the majority of the objective improvement is already accomplished in our optimizations and significant additional computational effect is required to reduce the objectives further. The optimizations are sufficiently converged to draw overall conclusions and to examine the design trade-offs in the resulting optimal design.

Next we compare the overall planform changes in the optimized designs. The chord, sweep and span variables, 7 in total, control the overall planform including the wing area. Figure 15 compares the new planforms with the original design. Despite the large amount of design flexibility the optimized designs have reasonable planforms. The MTOW optimization did not change the planform area but did increase the span from 58.7 m to 62.3 m and the aspect ratio from 9.0 to 10.1. The most striking feature, however, is the addition of a raked wingtip. While the planform design

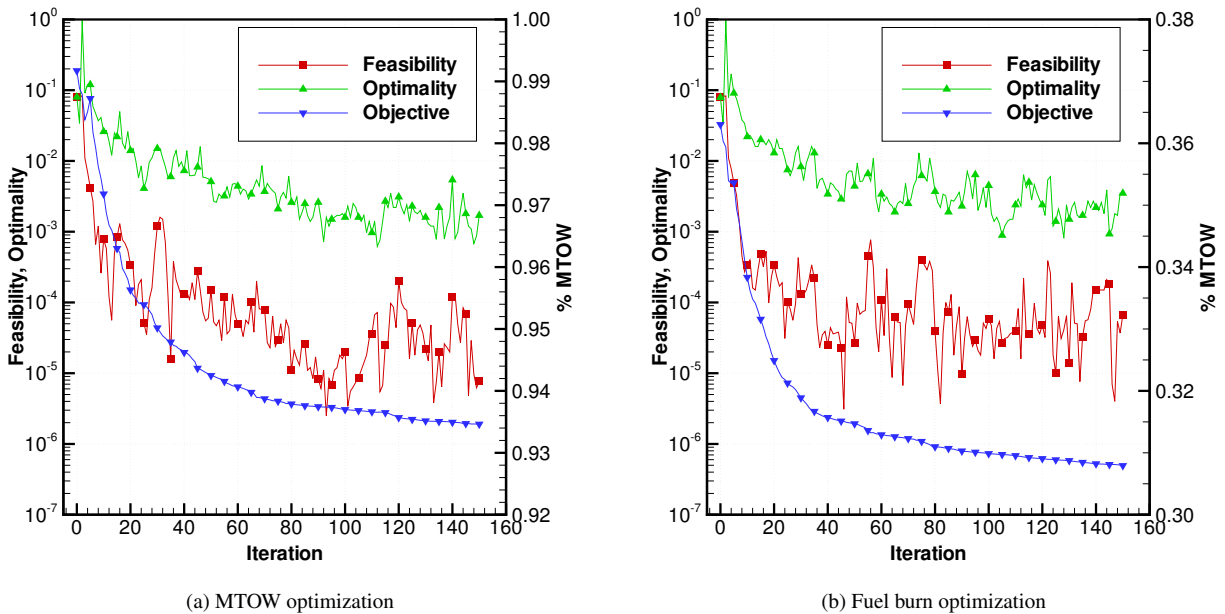


Figure 8: Optimization convergence histories for each optimization

variables are chosen in such a way to facilitate the exploration of this design feature by the optimization procedure, it only appears with the MTOW objective. In contrast, the fuel burn objective results in a significant increase in the wing span as well as wing area. For this case, the span has increased from 58.7 m to 72.7 m, the exposed planform area has increased by 16.2% and the new aspect ratio is 11.9. Looking closely at the tip, the leading edge is effectively straight right to the tip and there is a slight tapering of the chord near the tip. The leading sweep of the MTOW optimization remained essentially unchanged from the baseline except for the raked tip. The fuel burn minimization results in over 5° less sweep than the original design. The reduction in sweep is required to support the large span extension. The wave drag penalty for lowering the sweep angle is minimized through small shape modifications that can only be achieved with a large number of shape design variables. Additionally, the fuel burn optimum is thinner than either the original or MTOW optimum which further limits the negative aspects of lowering the sweep angle. These complex interactions can only be predicted with aerostructural optimization method.

Next we examine the evolution of the pressure drag (induce + wave drag), skin friction drag as well as the weight of the primary wing structure. The data is shown in Figure 9. Here we can clearly see the vastly different tradeoffs made by each optimization take to improve the chosen objective. The MTOW optimization slightly reduces the primary structural weight while simultaneously reducing the pressure drag. The skin friction drag increases slightly due primarily to the significant overall increases in the thickness to chord ratio. The fuel burn optimization on the other hand, requires a significantly heavier structure to support the large span increase. In this case, the increase in the skin friction drag is due primarily to the increase in the wing surface as the thickness to chord ratios decreased slightly. Despite the increased structural weight and increased skin friction drag, the total drag coefficient is lower for the fuel burn objective.

An examination of the span wise thickness to chord distribution and twist distributions can shed additional light on the design tradeoffs in each optimization. This information is plotted in Figure 10. We immediately see the vast difference between the optimized t/c ratios of the two optimized designs. For the MTOW optimized design, the t/c ratios have increased for the majority of the span, peaking at a maximum of 14.8% near the half span position. For the fuel burn optimization, an overall reduction in thickness to chord ratio is observed with the exception of slight increases near, and slightly outboard, of the yehudi break. It is now clear there are two contributing factors to the increased primary structural weight for the fuel burn optimized design: an increase in wing span resulting in higher bending moments and a decrease in the t/c distributing resulting in a less efficient structural design. The twist distributions plotted in Figure 10 are derived from the deformed, flying shape of the wing. The optimized twist distributions generally follow the original distribution with both designs lowering the amount of washout near the tip. The optimized twist distribution for the 2.5 g maneuver condition also sheds lights on the reason for the appearance of the raked wingtip for the MTOW optimization. Through a chord reduction and twist modification at the tip, the raked

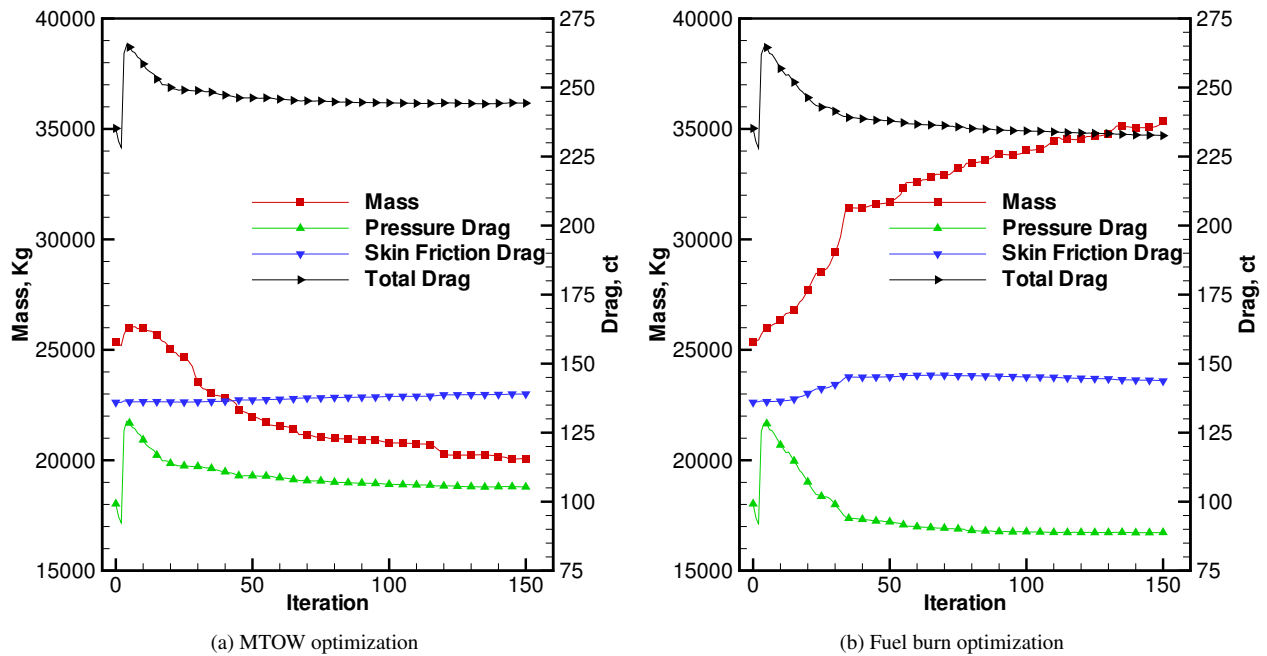


Figure 9: Weight and drag evolution during optimization (Cruise condition 1)

tip allows the maneuver condition to further reduce the loading near tip leading to a lighter structure. In the fuel burn case, even though the wing is heavier, the increased span leads to larger deflections and a higher overall flexibility. The same aeroelastic effect of reducing the tip load can be achieved naturally without raking the tip.

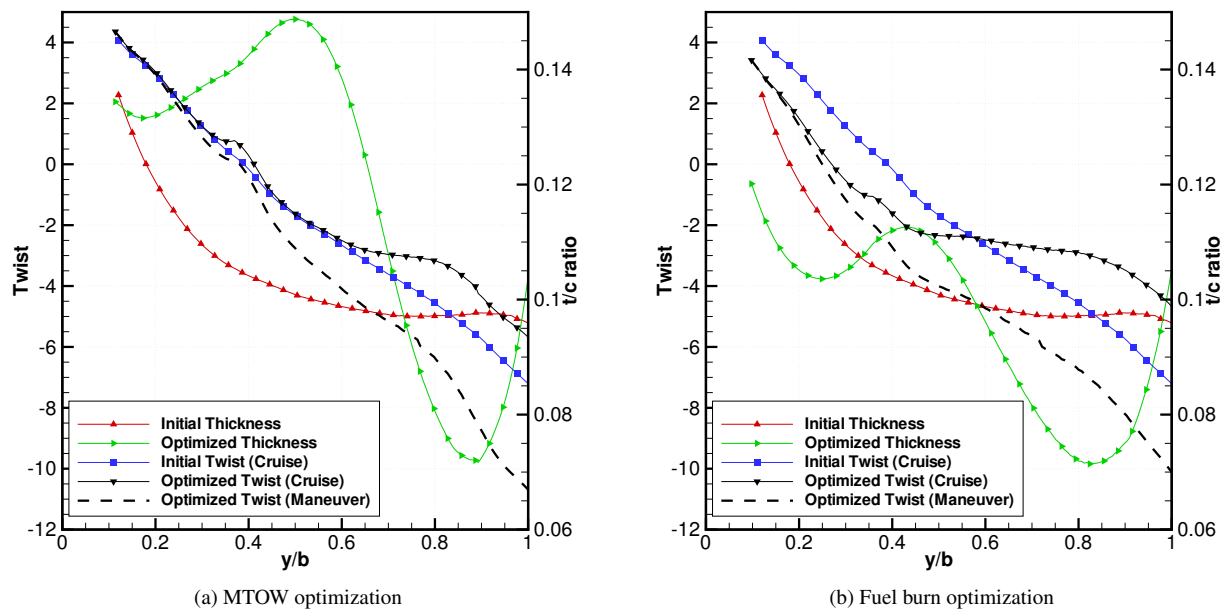


Figure 10: Twist and thickness to chord variation for the initial and optimized designs (Cruise condition 1)

Next we examine the span-wise lift distributions of the initial and optimized designs shown in Figure 11. The first observation is that for all designs the maneuver lift distributions have inboard shifted load distributions compared to the corresponding cruise lift distribution. This beneficial effect is caused by the aeroelastic induced washout near the tip, resulting in reduced loading. The cruise lift distributions for each optimization show more triangular load

distributions compared to the original or the elliptical reference. While the difference is slight, the fuel burn result has a lift distribution slightly closer to elliptical resulting in a slight reduction in induced drag. Comparing the lift distribution between the cruise and maneuver conditions, the difference is much more pronounced for the fuel burn result. As noted above, the lower tip load is a result of the more flexible structure and is achieved without a raked wing tip. Due to this flexibility, the maneuver load condition for the fuel burn optimum shows the highest root loading and lowest tip loading.

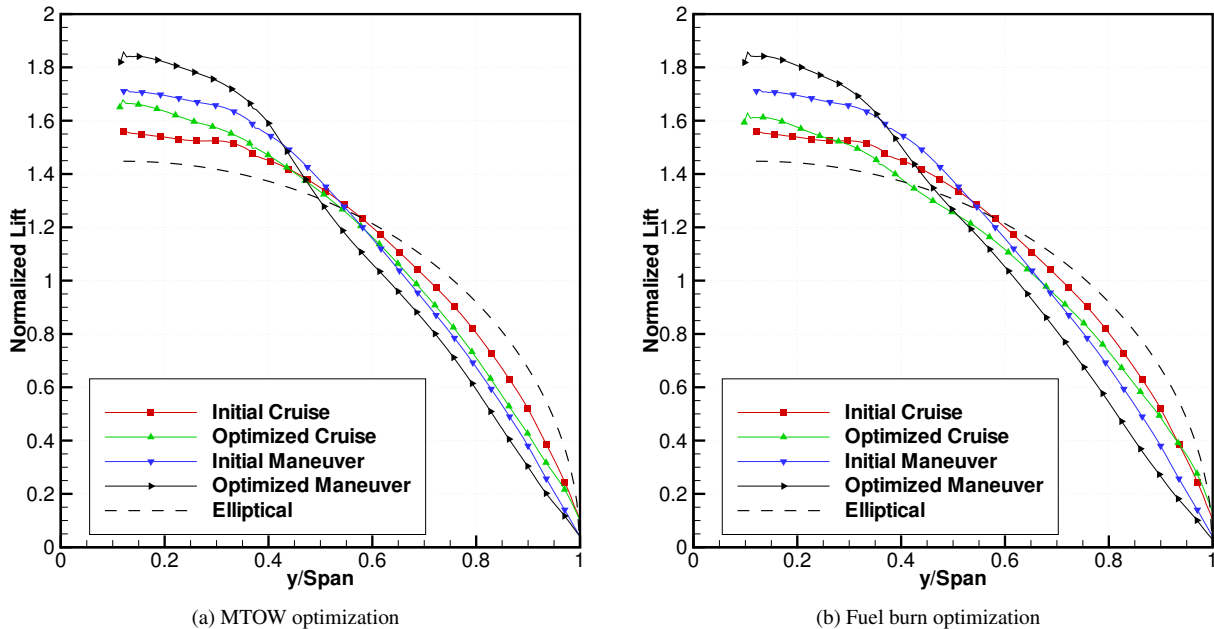


Figure 11: Span-wise lift distributions for the initial and optimized design for cruise condition 1 and maneuver condition 1

An examination of the chord-wise C_p distributions and airfoil shapes can shed light on the aerostructural tradeoffs the optimizer is required to make between the five cruise design points. Figure 13 shows the airfoil shape and C_p distribution for the initial and optimized designs for each operating condition at the 66% span location. Design points 1, 4 and 5 are all at the design cruise Mach number of 0.85 and all show a reduction in the shock strength on the wing upper surface lowering the wave drag. Design point 1, at a lower Mach number of 0.84, shows a distinct, weak double shock structure on the upper wing surface. The higher mach number case of 0.86, (design point 2) shows a slightly reduced shock strength compared to the original. If we closely examine the airfoil cross section for the MTOW optimization, we observe a significant thickening of the aft portion of the airfoil. This is primarily to allow a thicker, more efficient load carrying structure which helps explain the mass reduction achieved for the MTOW optimization despite an increase in wingspan. In fact, at this span-wise location, the vertical displacement of the optimized design is lower. This aft thickening is present on the fuel burn optimization but to a lesser extent. This thickening effect leads, however leads to a much more rapid pressure recovery than the original which increases the risk of flow separation which should be investigated with a RANS flow solution. A comparison of the deflections of the initial and optimized design as well as the jig shape is given in Figure 12.

Now we examine the changes in the structural design from the two aerostructural optimizations. Figure 14 shows the initial design and the two optimized configurations. Recall the initial design was first optimized with fixed aerostructural loads and this mass minimization resulted in a failure distribution (similar to Figure 14b or Figure 14f) that just meet the KS failure constraints. However, when this design is used in the multidisciplinary optimization, we observe lower stress values. This highlights the need for multidisciplinary optimization: Even though we performed a structural optimization, we would have to continue the procedure of generating loads from the previous design and re-optimizing to generate a design that just meets the multidisciplinary constraint.

Each optimization resulted in similar failure distributions. The plots in Figure 14 show the failure load for maneuver condition 1, the symmetric 2.5 g pull up maneuver. The upper skin and stringers are critical for this load case, but not the lower surface; these are dimensioned by the maneuver load 2, the 1.3 g gust load. The overall distribution of material for the initial and two optimized designs is similar. The largest skin thickness occurs through the mid




Figure 12: Front view showing aerostructural deflections for the 2.5g maneuver condition (left) and cruise design point 1 (right)

section of the spar box where the thickness is greatest. The MTOW optimum shows locally increased thickness need the yehudi break of the wing even with the large increase in t/c ratios we saw earlier. In order to support loading from the span extension, the fuel burn results shown significantly increased thicknesses over the majority of the lower skin and stringers.

Finally, Table 12 lists the key results from the optimizations. Drag reduction for all operating conditions for both optimizations is achieved with the higher Mach number point (2) and the higher lift condition (5) resulting in the largest improvements.

For the fuel burn optimization the drag lowered by 30%. As with the MTOW optimization, there is a reduction in wave drag, but there is a much greater reduction in the induced drag due to the much higher aspect ratio. In fact, the MTOW for the fuel burn optimization remained essentially unchanged and lift cruise lift constraints, imposed at the midpoint of cruise are actually higher for the fuel burn case.

A weight breakdown of the primary structure sheds additional light on where the weight is saved for MTOW optimization. The largest weight reduction is from the top and bottom skins and the lower stringers. For the lower skin and stringers, maneuver load case 2 is critical. From Figure 11, even the 1 g cruise conditions have more heavily loaded inboard lift distribution than the initial design. Since the 1.3 g lift distribution is similar to the cruise condition, we can conclude the slight induced drag penalty from a more triangular lift distribution is offset by the weight reduction in the lower skin and stringers. For the fuel burn optimization, we see weight increases across essentially all components with the largest increases from the skins and rear spar. Given the reduced t/c ratio and increased span, these increases are expected.

V. Conclusions

A framework for the analysis and optimization of high-fidelity aerostructural systems was presented. Two methods for solving the nonlinear aerostructural systems were presented: a block Gauss–Seidel method with Aitken acceleration and a fully coupled Newton–Krylov approach. Both methods perform well on our problem of interest, with the Newton–Krylov method typically requiring 10% less computational time than the traditional segregated, block Gauss–Seidel method. A typical aerostructural solution with 2 million CFD cells and 300 000 structural degrees of freedom can be obtained in under 3 minutes using 52 aerodynamic and 4 structural processors.

A coupled adjoint for the aerostructural equations was implemented that significantly improved on previous work. By eliminating the inaccuracy and computational expense of finite-differencing for the off-diagonal terms, we are able to solve much larger coupled adjoint systems with over 80 million degrees of freedom very efficiently. A traditional linear block Gauss–Seidel method was compared with a new fully coupled Krylov method implementation on a 2 million cell CFD mesh and a 300 000 DOF structural mesh. We showed that the fully coupled solution method is approximately 30% faster than the linear block Gauss–Seidel method with similar memory requirements.

A parallel scaling study demonstrated the techniques presented in the paper can be scaled to CFD problems with over 16 millions cells and CSM discretizations with over one million degrees of freedom and approximately 100 000 coupling variables. A verification of the accuracy of the coupled adjoint method with the complex-step derivative approximation showed excellent agreement with relative errors of $\mathcal{O}(10^{-5})$.

Using the optimization framework, aerostructural optimizations were performed on the Common Research Model geometry with a structure representative of a modern airliner wing. Multi-point optimizations with 5 cruise conditions and 2 maneuver conditions were performed with a 2 million cell CFD mesh and 300 000 DOF structural mesh. The optimization problems used 476 design variables and 57 constraints and required 36 hours of wall time on 435 processors.

Two objective functions, MTOW and fuel burn, were considered which provided insights into the aerostructural tradeoffs the analysis can capture. A sensitivity analysis of each objective function showed the fuel burn objective should result in a more “aerodynamic” solution with lower drag than the MTOW objective. The optimized results show that this is indeed the case: the fuel burn optimization reduced the cruise segment fuel consumption by 11.2%, while the MTOW optimization reduced the fuel burn by 6.6%. Conversely, the MTOW minimization resulted in 4.2% reduction in MTOW while the MTOW for the fuel burn optimization remained effectively unchanged. The optimization results demonstrate the utility of the framework for aerostructural design optimization of full configurations with hundreds of design variables and real world constraints.

VI. Future Work

During the course of this research we have identified a number of areas that warrant further investigation. The most significant improvement would be to replace the Euler equations with the Reynolds averaged Navier–Stokes (RANS) equations for the aerodynamic solution. The framework can easily be extended to include this addition and work is currently underway to improve our current RANS solution technique and to implement the additional partial derivative terms required for the coupled adjoint. The RANS equations would not only yield a more accurate prediction of the skin friction drag, but would improve the prediction of the shock strength and location.

The structural analysis in this work only considered a linear structural response. The large deflections observed for the 2.5 g maneuver condition calls this assumption into question. TACS is capable of analyzing a geometrically nonlinear structural response and this can be incorporated into the framework with relatively easily. We also would like to consider composite wing box structures in the future, a capability already included in TACS. Composite structures are playing an increasing important role in commercial aircraft design and with this capability integrated into the MACH framework, we would like to investigate the aerostructural optimization of highly flexible composite structures.

As was noted in Section II.5, more attention needs to be given to the startup and convergence behavior of the Newton–Krylov solution technique. We intend to investigate other continuation strategies, such as pseudo-transient continuation, to improve the convergence and startup behavior of the method.

Finally, as we begin to examine more operating conditions, it becomes increasingly important to decide how to combine the performance parameters from each analysis into a composite objective function that is realistic. A weighted average can be used to assign more importance to certain operating conditions, but how these weights are selected and their effect on the optimum design is not trivial. We are currently investigating the use of surrogate models to approximate the aerostructural performance in the cruise regime to allow more complex mission analysis and incorporate actual aircraft usage patterns into the objective formulation.

VII. Acknowledgments


The authors are grateful for the funding provided by the Natural Sciences and Engineering Research Council. Computations were performed on the GPC supercomputer at the SciNet HPC Consortium. SciNet is funded by: the Canada Foundation for Innovation under the auspices of Compute Canada; the Government of Ontario; Ontario Research Fund – Research Excellence; and the University of Toronto.

References


- [1] Maute, K., Nikbay, M., and Farhat, C., “Coupled analytical sensitivity analysis and optimization of three-dimensional nonlinear aeroelastic systems,” *AIAA Journal*, Vol. 39, No. 11, 2001, pp. 2051–2061.
- [2] Maman, N. and Farhat, C., “Matching fluid and structure meshes for aeroelastic computations: a parallel approach,” *Computers and Structures*, Vol. 54, No. 4, 1995, pp. 779–785.
- [3] Farhat, C., Lesoinne, M., and LeTallec, P., “Load and motion transfer algorithms for fluid/structure interaction problems with non-matching discrete interfaces: Momentum and energy conservation, optimal discretization and application to aeroelasticity,” *Computer methods in applied mechanics and engineering*, 1996, pp. 95–115.
- [4] Martins, J. R. R. A., Alonso, J. J., and Reuther, J. J., “Complete Configuration Aero-Structural Optimization Using a Coupled Sensitivity Analysis Method,” *Proceedings of the 9th AIAA/ISSMO Symposium on Multidisciplinary Analysis and Optimization*, Atlanta, GA, September 2002.
- [5] Martins, J. R. R. A., Martins, J. R. R. A., Alonso, J. J., Alonso, J. J., Reuther, J. J., and Reuther, J. J., “High-Fidelity Aero-Structural Design Optimization of a Supersonic Business Jet,” *Journal of Aircraft*, Vol. 41, No. 3, 2004, pp. 523–530. doi:10.2514/1.11478.

- [6] Martins, J. R. R. A., Alonso, J. J., and Reuther, J. J., “A Coupled-Adjoint Sensitivity Analysis Method for High-Fidelity Aero-Structural Design,” *Optimization and Engineering*, Vol. 6, No. 1, March 2005, pp. 33–62. doi:10.1023/B:OPTE.0000048536.47956.62.
- [7] Maute, K., Nikbay, M., and Farhat, C., “Sensitivity analysis and design optimization of three-dimensional non-linear aeroelastic systems by the adjoint method,” *International Journal for Numerical Methods in Engineering*, Vol. 56, 2003, pp. 911–933.
- [8] Barcelos, M., Bavestrello, H., and Maute, K., “A Schur-Newton-Krylov solver for steady-state aeroelastic analysis and design sensitivity analysis,” *Computer Methods in Applied Mechanics and Engineering*, Vol. 195, No. 17-18, 2006, pp. 2050 – 2069. doi:10.1016/j.cma.2004.09.013.
- [9] Barcelos, M. and Maute, K., “Aeroelastic design optimization for laminar and turbulent flows,” *Computer Methods in Applied Mechanics and Engineering*, Vol. 197, No. 19-20, 2008, pp. 1813 – 1832. doi:10.1016/j.cma.2007.03.009.
- [10] Alonso, J. J., LeGresley, P., van der Weide, E., Martins, J. R. R. A., and Reuther, J. J., “pyMDO: A Framework for High-Fidelity Multi-Disciplinary Optimization,” *Proceedings of the 10th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*, Albany, NY, September 2004, AIAA 2004-4480.
- [11] Hascoët, L. and Pascual, V., “TAPENADE 2.1 User’s Guide,” Technical report 300, INRIA, 2004.
- [12] van der Weide, E., Kalitzin, G., Schluter, J., and Alonso, J. J., “Unsteady Turbomachinery Computations Using Massively Parallel Platforms,” *Proceedings of the 44th AIAA Aerospace Sciences Meeting and Exhibit*, Reno, NV, 2006, AIAA 2006-0421.
- [13] Balay, S., Gropp, W. D., McInnes, L. C., and Smith, B. F., “Efficient Management of Parallelism in Object Oriented Numerical Software Libraries,” *Modern Software Tools in Scientific Computing*, edited by E. Arge, A. M. Bruaset, and H. P. Langtangen, Birkhäuser Press, 1997, pp. 163–202.
- [14] Kennedy, G. J. and Martins, J. R. R. A., “Parallel Solution Methods for Aerostructural Analysis and Design Optimization,” *Proceedings of the 13th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*, Fort Worth, Texas, September 2010.
- [15] Brown, S., “Displacement extrapolation for CFD+CSM aeroelastic analysis,” AIAA Paper 97-1090, 1997.
- [16] Hicken, J. and Zingg, D., “Aerodynamic Optimization Algorithm with Integrated Geometry Parameterization and Mesh Movement,” *AIAA Journal*, Vol. 48, No. 2, February 2009, pp. 400–413.
- [17] Kenway, G. K. W. and Joaquim R.R.A. Martins, G. J. K., “A CAD-Free Approach to High-Fidelity Aerostructural Optimization,” *13th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*, Fort Worth, TX, September 2010, AIAA-2010-9231.
- [18] Martins, J. R. R. A. and Lambe, A. B., “Multidisciplinary design optimization: Survey of architectures,” *AIAA Journal*, 2012, (Submitted).
- [19] Tedford, N. P. and Martins, J. R. R. A., “Benchmarking Multidisciplinary Design Optimization Algorithms,” *Optimization and Engineering*, Vol. 11, No. 1, February 2010, pp. 159–183. doi:10.1007/s11081-009-9082-6.
- [20] Irons, B. M. and Tuck, R. C., “A Version of the Aitken Accelerator for Computer Iteration,” *International Journal for Numerical Methods in Engineering*, Vol. 1, 1969, pp. 275–277.
- [21] Saad, Y. and Schultz, M. H., “GMRES: A Generalized Minimal Residual Algorithm for Solving Nonsymmetric Linear Systems,” Vol. 7, No. 3, 1986, pp. 856–869. doi:10.1137/0907058.
- [22] Eisenstat, S. C., Walker, H. F., Eisenstatt, S. C., Homer, and Walker, F., “Choosing the Forcing Terms in an Inexact Newton Method,” *SIAM J. Sci. Comput.*, Vol. 17, 1994, pp. 16–32.
- [23] Hopkins, E. J., “Charts for predicting turbulent skin friction from the Van Driest Method (II),” Tech. rep., NASA Ames Research Center, 1972, TN D-6945.
- [24] Mason, W., “Program FRICTION,” Website, http://www.dept.aoe.vt.edu/~mason/Mason_f/MRsoft.html.
- [25] Peterson, P., Martins, J. R. R. A., and Alonso, J. J., “Fortran to Python Interface Generator with an Application to Aerospace Engineering,” *Proceedings of the 9th International Python Conference, Long Beach, CA*, 2001.
- [26] Martins, J. R. R. A. and Hwang, J. T., “Review of Options for Computing Derivatives of Multidisciplinary Systems,” *53rd AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference*, Honolulu, HI, April 2012.
- [27] Mader, C. A., Martins, J. R. R. A., Alonso, J. J., and van der Weide, E., “ADjoint: An Approach for the Rapid Development of Discrete Adjoint Solvers,” *AIAA Journal*, Vol. 46, No. 4, April 2008, pp. 863–873. doi:10.2514/1.29123.
- [28] Goldfarb, D. and Toint, P. L., “Optimal Estimation of Jacobian and Hessian Matrices That Arise in Finite Difference Calculations,” *Mathematics of Computation*, Vol. 43, 1984, pp. 69–88.
- [29] Li, X. S. and Demmel, J. W., “SuperLU DIST: A scalable distributed-memory sparse direct solver for unsymmetric linear system,” *ACM Trans. Mathematical Software*, Vol. 29, No. 2, 2003, pp. 110–140.
- [30] Raspanti, C., Bandoni, J., and Biegler, L., “New strategies for flexibility analysis and design under uncertainty,” *Computers and Chemical Engineering*, Vol. 24, 2000, pp. 2193–2209.


- [31] Wrenn, G., "An indirect method for numerical optimization using the Kreisselmeier-Steinhauser function," NASA Technical Report CR-4220, 1989.
- [32] Poon, N. M. K. and Martins, J. R. R. A., "An Adaptive Approach to Constraint Aggregation Using Adjoint Sensitivity Analysis," *Structures and Multidisciplinary Optimization*, Vol. 30, No. 1, 2007, pp. 61–73. doi:[10.1007/s00158-006-0061-7](https://doi.org/10.1007/s00158-006-0061-7).
- [33] Vassberg, J. C., DeHaan, M. A., Rivers, S. M., and Wahls, R. A., "Development of a Common Research Model for Applied CFD Validation Studies," *26th AIAA Applied Aerodynamics Conference*, AIAA, Honolulu, Hawaii, August 2008.
- [34] "4th AIAA CFD Drag Prediction Workshop," Internet, June 2009, <http://aaac.larc.nasa.gov/tsab/cfdlarc/aiaa-dpw/>.
- [35] Boeing Commercial Airplanes, "777-200/300 Airplane Characteristics for Airport Planning," Internet, July 1998, D6-58329.
- [36] Martins, J. R. R. A., Sturdza, P., and Alonso, J. J., "The Complex-Step Derivative Approximation," *ACM Transactions on Mathematical Software*, Vol. 29, No. 3, 2003, pp. 245–262. doi:[10.1145/838250.838251](https://doi.org/10.1145/838250.838251).
- [37] Smith, B., Bjørstad, P., and Gropp, W., *Domain Decomposition: Parallel Multilevel Methods for Elliptic Partial Differential Equations*, Cambridge University Press, 2004.
- [38] Loken, C., Gruner, D., Groer, L., Peltier, R., Bunn, N., Craig, M., Henriques, T., Dempsey, J., Yu, C.-H., Chen, J., Dursi, L. J., Chong, J., Northrup, S., Pinto, J., Knecht, N., and Zon, R. V., "SciNet: Lessons Learned from Building a Power-efficient Top-20 System and Data Centre," *Journal of Physics: Conference Series*, Vol. 256, No. 1, 2010, pp. 012026.
- [39] Nemeec, M., Zingg, D. W., and Pulliam, T. H., "Multipoint and Multi-Objective Aerodynamic Shape Optimization," *AIAA Journal*, Vol. 42, No. 6, June 2004, pp. 1057–1065.
- [40] Mader, C. A. and Martins, J. R. R. A., "Optimal Flying Wings: A Numerical Optimization Study," *53rd AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference*, Honolulu, HI, April 2012, AIAA 2012-.
- [41] Gill, P. E., Murray, W., and Saunders, M. A., "SNOPT: An SQP Algorithm for Large-Scale Constrained Optimization," *SIAM Review*, Vol. 47, No. 1, 2005, pp. 99–131. doi:[10.1137/S0036144504446096](https://doi.org/10.1137/S0036144504446096).




(a) MTOW optimization, point 1




(b) Fuel burn optimization, point 1




(c) MTOW optimization, point 2




(d) Fuel burn optimization, point 2




(e) MTOW optimization, point 3




(f) Fuel burn optimization, point 3




(g) MTOW optimization, point 4



(h) Fuel burn optimization, point 4




(i) MTOW optimization, point 5




(j) Fuel burn optimization, point 5


Figure 13: Cross section and C_p contours for each operating condition at 66% span. Initial design: dashed lines; optimized design: solid lines.




(a) Thickness distribution for MTOW optimization



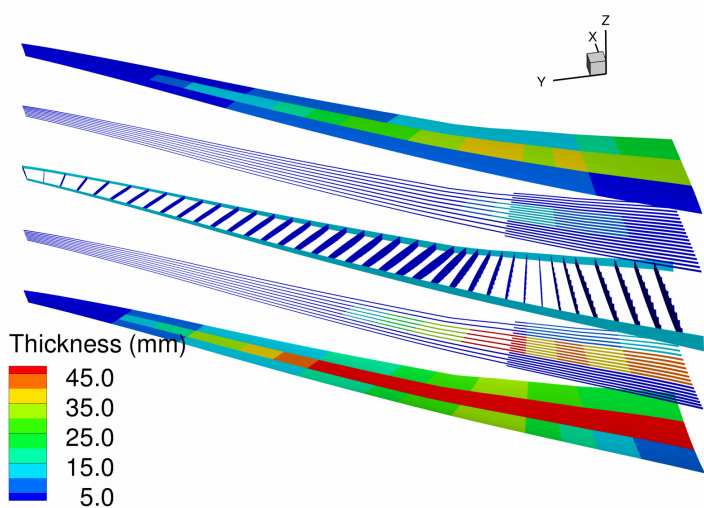
(b) Failure distribution for MTOW optimization, maneuver condition 1




(c) Thickness distribution for initial design



(d) Failure distribution for initial design, maneuver condition 1

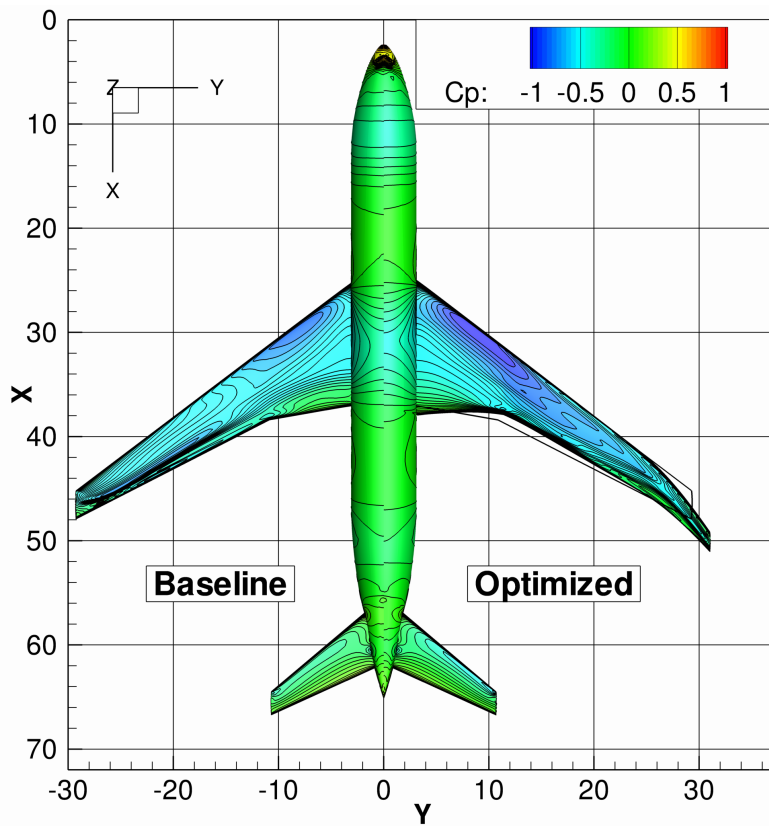


(e) Thickness distribution for fuel burn optimization




(f) Failure distribution for fuel burn optimization, maneuver condition 1

Figure 14: Initial and optimized thickness (left column) and failure parameter (right column) distributions over the wing box



(a) MTOW optimization



(b) Fuel burn optimization

Figure 15: Planform view with C_p contours for both optimum designs for cruise condition 1. The initial design, on left, is included for reference

Table 12: Optimization results summary

		Initial	MOTW optimization		Fuel burn optimization	
		Value	Value	% Change	Value	% Change
Induced + wave drag (ct)	Cruise 1	134.0	105.3	-21.4%	88.9	-33.7%
	Cruise 2	133.1	112.6	-15.4%	93.2	-30.0%
	Cruise 3	143.4	111.5	-22.2%	93.8	-34.6%
	Cruise 4	122.7	99.4	-19.0%	83.7	-31.8%
	Cruise 5	147.7	114.3	-22.2%	95.9	-35.1%
	Average	136.0	108.6	-20.2%	91.1	-33.1%
Friction drag (ct)		136.0	139.0	2.2%	143.7	5.7%
Mass breakdown (kg)	Bottom skin	12 653	9 192	-27.4%	18 583	46.9%
	Bottom stringers	2 055	1 687	-17.9%	2 539	23.6%
	Top skin	7 519	5 424	-27.9%	10 545	40.3%
	Top stringers	557	691	24.0%	634	13.8%
	Rib stiffeners	45	56	24.5%	51	14.4%
	Ribs	581	578	-0.5%	578	-0.5%
	Front spar	1 235	1 307	5.8%	1 387	12.4%
	Rear spar	723	1 145	58.3%	1 049	45.1%
	Total	25 368	20 078	-20.8%	35 368	39.4%
Final cruise weight (kgf)		187 071	181779	-2.8%	198 367	6.0%
Initial cruise weight (kgf)	Cruise 1	289 203	276 875	-4.3%	289 148	0.0%
	Cruise 2	287 321	278 618	-3.0%	289 886	0.9%
	Cruise 3	295 213	280 218	-5.1%	291 991	-1.1%
	Cruise 4	290 405	278 340	-4.2%	290 729	0.1%
	Cruise 5	288 858	276 411	-4.3%	288 201	-0.2%
	Average	290 200	278 098	-4.2%	289 991	-0.1%
Fuel burn (kgf)	Cruise 1	102 132	95 097	-6.9%	90 782	-11.1%
	Cruise 2	100 250	96 839	-3.4%	91 520	-8.7%
	Cruise 3	108 141	98 440	-9.0%	93 625	-13.4%
	Cruise 4	103 334	96 561	-6.6%	92 362	-10.6%
	Cruise 5	101 789	96 320	-7.0%	89 834	-11.7%
	Average	103 129	96 320	-6.6%	91 625	-11.2%
Span (m)		58.6	62.1	6.0%	72.3	23.4%
Aspect ratio		9.0	10.1	12.2%	11.9	32.2%
Leading edge sweep (°)		37.4	37.6	0.1%	32.0	-14.4%
Reference wing area (m²)		383.7	383.7	0.0%	437.9	14.1%
Exposed wing area (m²)		335.0	335.0	0.0%	389.2	16.2%
Wing volume (m³)		46.0	53.0	15.2%	48.7	5.9%