Aerodynamic Shape Optimization of Benchmark Problems Using Jetstream

Christopher Lee*, David Koo†, Karla Telidetzki‡, Howard Buckley§, Hugo Gagnon¶, and David W. Zingg‖

Institute for Aerospace Studies, University of Toronto

4925 Dufferin St., Toronto, Ontario, M3H 5T6, Canada

This work presents results from the application of an aerodynamic shape optimization code, Jetstream, to a suite of benchmark cases defined by the Aerodynamic Design Optimization Discussion Group. Geometry parameterization and mesh movement are integrated by fitting the multi-block structured grids with B-spline volumes and performing mesh movement based on a linear elastic model applied to the control points. Geometry control is achieved through two different approaches. Either the B-spline surface control points are taken as the design variables for optimization, or alternatively, the surface control points are embedded within free-form deformation (FFD) B-spline volumes, and the FFD control points are taken as the design variables. Spatial discretization of the Euler or Reynolds-averaged Navier-Stokes equations is performed using summation-by-parts operators with simultaneous approximation terms at boundaries and block interfaces. The governing equations are solved iteratively using a parallel Newton-Krylov-Schur algorithm. The discrete-adjoint method is used to calculate the gradients supplied to a sequential quadratic programming optimization algorithm. The first optimization problem studied is the drag minimization of a modified NACA 0012 airfoil at zero angle of attack in inviscid, transonic flow, with a minimum thickness constraint set to the initial thickness. The shock is weakened and moved downstream, reducing drag by 91%. The second problem is the lift-constrained drag minimization of the RAE 2822 airfoil in viscous, transonic flow. The shock is eliminated and drag is reduced by 48%. Both two-dimensional cases exhibit optimization convergence difficulties due to the presence of nonunique flow solutions. The third problem is the twist optimization for minimum induced drag at fixed lift of a rectangular wing in subsonic, inviscid flow. A span efficiency factor very close to unity and a near elliptical lift distribution are achieved. The final problem includes single-point and multi-point lift-constrained drag minimizations of the Common Research Model wing in transonic, viscous flow. Significant shape changes and performance improvements are achieved in all cases. Finally, two additional optimization problems are presented that demonstrate the capabilities of Jetstream and could be suitable additions to the Discussion Group problem suite. The first is a wing-fuselage-tail optimization with a prescribed spanwise load distribution on the wing. The second is an optimization of a box-wing geometry.

I. Introduction

The dual forces of growing concern over the negative impact of carbon emissions in the environment and the rise of jet fuel prices pressure aircraft manufacturers to prioritize minimizing fuel burn when designing new aircraft. In 2012, commercial flights transported close to 3 billion passengers around the world.[1] In

*MASc Candidate, Student Member AIAA.

†MASc Candidate

‡MASc Candidate (currently Methods Agent, Bombardier Aerospace)

§Research Engineer, Member AIAA.

¶PhD Candidate, Student Member AIAA.

‖Professor and Director, Tier 1 Canada Research Chair in Computational Aerodynamics and Environmentally Friendly Aircraft Design, J. Armand Bombardier Foundation Chair in Aerospace Flight, Associate Fellow AIAA.

the process, 72 billion gallons of jet fuel were consumed, releasing 682 million tonnes of CO_2 emissions into the atmosphere.[2] It is expected that the annual number of passengers will double by 2030,[1] increasing the potential impact on the environment. In efforts to minimize the aircraft industry's environmental footprint, the industry has committed to "improving fuel efficiency an average of 1.5% annually to 2020, capping net emissions through carbon-neutral growth from 2020, [and] cutting net emissions in half by 2050, compared with 2005."[3] At the same time, fuel costs have more than doubled between 2004 and 2013, and can now account for 31% of an airline's operating costs.[3] With fuel prices rising, operation of fuel efficient aircraft is key to ensuring the profitability of airlines. In short, continual improvements in fuel efficiency are required to ensure both environmental and economic sustainability of air transport.

One key design focus for aircraft manufacturers in meeting this concern is drag minimization in aerodynamic design using computational fluid dynamics (CFD). On its own, a CFD solver is a tool capable of analyzing only a specific design, but in recent years the rapid development of computing power has enabled the feasible use of computational design tools which have not only sped up, but drastically altered the design process. By coupling a CFD solver with an optimization algorithm and a geometry parameterization tool, designers are able to perform aerodynamic shape optimization, robustly exploring a design space at a fraction of the financial and time cost that would be needed for an experimental cut and try approach, or to manually alter models for numerical analysis. These tools allow not only for more robust fine tuning of existing designs, but also the exploration of unconventional configurations which may provide more dramatic aerodynamic improvements.

Forums for comparison of CFD algorithms that enable different research groups to validate their codes are already well established, such as the Drag Prediction Workshops.[4] In the same spirit, researchers in the aerodynamic design optimization community have initiated the AIAA Aerodynamic Design Optimization Discussion Group (ADODG)[a]. The ADODG has defined a series of benchmark optimization problems, allowing research groups from industry and academia to test and compare their codes under a variety of problems. A range of flow conditions and geometric flexibility are considered. The first meeting for the discussion group was in 2014, and this year, researchers reconvene with updated results for the cases, in light of lessons learned from last year.

The four benchmark cases are constrained drag minimization problems. The first case is the sectional optimization of a modified NACA 0012 airfoil at zero angle of attack in inviscid, transonic flow, with a minimum thickness constraint set to the initial thickness. The second case is the sectional optimization of an RAE 2822 airfoil in viscous, transonic flow, subject to lift, pitching moment, and area constraints. The third case is the twist optimization of a rectangular wing in subsonic, inviscid flow, subject to a lift constraint. The final case is the sectional and twist optimization of the Common Research Model (CRM) wing in transonic, viscous flow, subject to lift, pitching moment, volume, and thickness constraints. This final case includes a single-point optimization and multi-point optimizations at varying lift coefficients and Mach numbers. This paper presents the updated results obtained for these cases since last year's meeting.[5]

The remainder of this paper is outlined as follows: Section II summarizes the algorithms employed in the aerodynamic design optimization framework Jetstream. Section III presents the results obtained for the four benchmark problems. Section IV proposes two potential new cases to be added to the benchmark suite, along with results. Section V outlines conclusions and future work.

II. Methodology

A. Integrated Geometry Parameterization, Control, and Mesh Movement

B-Spline Surface Geometry Parameterization and Control

Jetstream uses a cubic B-spline surface parameterization that can accurately capture an initial geometry while providing good geometric flexibility with a modest number of design variables.[6] Each block in the multi-block mesh is fitted with a cubic B-spline volume with a specified number of control points, and the control points defining the geometry's surface are taken as the design variables.

The fitting procedure is described as follows. The parametric values of the grid nodes $G = \{x_{q,r,s} | q = 1, ...L_q, r = 1, ..., L_r, s = 1, ..., L_s\}$ are located. For example, parameter $\xi_{q,r,s}$ is calculated along the grid line

[a]https://info.aiaa.org/tac/ASG/APATC/AeroDesignOpt-DG/default.aspx

of constant $r = r_0$ and $s = s_0$ as

$$\xi_{1,r_0,s_0} = 0 \tag{1a}$$

$$\xi_{q,r_0,s_0} = \frac{1}{\Psi}\sum_{t=1}^{q-1} ||\mathbf{x}_{t+1,r_0,s_0} - \mathbf{x}_{t,r_0,s_0}||, \quad q = 2, ..., L_q, \tag{1b}$$

where the normalization factor of total arc length is given by

$$\Psi = \sum_{t=1}^{L_q-1} ||\mathbf{x}_{t+1,r_0,s_0} - \mathbf{x}_{t,r_0,s_0}||. \tag{2}$$

Next, the knot vectors are determined. To allow for a more accurate mapping, the knots are generalized to be spatially varying in parametric space. Bilinear knots have been chosen for simplicity:

$$T_i(\eta,\zeta) = T_{i,(0,0)}(1-\eta)(1-\zeta) + T_{i,(1,0)}\eta(1-\zeta) + T_{i,(0,1)}(1-\eta)\zeta + T_{i,(1,1)}\eta\zeta. \tag{3}$$

For brevity, the edge knot equations are omitted, but are described in Hicken.[7] Finally, the control-point coordinates are determined by solving a least-squares problem to best fit the grid to the initial geometry.

Free-Form Deformation Geometry Control

Free-form deformation (FFD) is the second geometry control method used in this paper. Conceptually, FFD can be visualized by imagining the geometry as a flexible object, enclosing the geometry in a larger volume of flexible material, and indirectly deforming the geometry by deforming the enclosing volume. In traditional aerodynamic design optimization practice, the embedded objects are the surface grid nodes. In Jetstream, however, the embedded objects are taken as the B-spline surface control points defining the geometry, and the FFD volume is a cubic B-spline volume.[8] This maintains an analytical definition of the geometry and allows the mesh movement, described later, to be performed in the same way as with B-spline surface control. The FFD volume is created using a geometry generation tool called GENAIR.[9]

Numerically, FFD is executed using two functions. The first, $F^{-1}(\mathbf{t}) = \boldsymbol{\xi}$, is an embedding function evaluated only once and is a mapping from world space $\mathbf{t}$ to parametric space $\boldsymbol{\xi}$. In this case, the surface control points are mapped to the parametric space of the FFD volume. The second function, $\tilde{F}(\boldsymbol{\xi}) = \tilde{\boldsymbol{t}}$, is a deformation function which algebraically re-evaluates the coordinates of every embedded surface control point once the FFD volume lattice points $\{\mathbf{B}_{i,j,k}\}$ have been adjusted by the optimization.

While B-spline surface control couples the design variables with the geometry parameterization, FFD decouples the two, parameterizing deformations rather than the geometry itself. So while the geometric design variables in the surface-based parameterization approach are the surface control points, the geometric design variables in the FFD approach are a set of the FFD volume control points $\mathbf{v}_{\text{geo}} \in \{\mathbf{B}_{i,j,k}\}$. Gagnon and Zingg[8] describe the deformation process as a two-level approach. The first level involves the control points defining the FFD volume. The second level involves the control points defining the geometry.

Linear-Elasticity Mesh Movement

Aerodynamic design optimization algorithms require some method to update the computational mesh once the geometry has been modified. Mesh regeneration is often too expensive and difficult to automate, so mesh movement methods are often preferred. The mesh movement method employed in Jetstream is based on a linear-elasticity model.[10] While such models can be expensive if applied directly to the computational mesh, Jetstream makes use of the fact that the fitted B-spline mesh acts as a control mesh providing a coarser approximation to the computational mesh. By applying the linear-elastic model to the control mesh rather than the computational mesh, the mesh movement becomes much cheaper to compute while still maintaining high mesh quality.[6]

The control mesh can be visualized as a solid that is elastically deformed in response to the displaced surface control points. Hexahedral cells defined by adjacent control points are assigned greater stiffness in inverse proportion to their volume to maintain the quality of the initial mesh. The system solved is defined by:

$$\boldsymbol{\mathcal{M}} = \mathbf{K}(\mathbf{b} - \mathbf{b}^{(0)}) - \mathbf{f} = 0, \tag{4}$$

where $\boldsymbol{\mathcal{M}}$ are the mesh residuals, $\mathbf{K}$ is the stiffness matrix, $\mathbf{b}$ and $\mathbf{b}^{(0)}$ are the updated and initial control point coordinate column vectors, respectively, and $\mathbf{f}$ is the force vector implicitly determined from the displaced surface control points. The mesh movement can be performed in increments to improve mesh quality for large shape changes. Since the original mesh fitting provides the parametric values of the grid nodes, algebraic recomputation of their coordinates in physical space is quick to perform.

B. Flow Solver

The flow solver in Jetstream is a three-dimensional multi-block structured finite-difference solver. The parallel implicit solver uses a Newton-Krylov-Schur method and is capable of solving the Euler or Reynolds-averaged Navier-Stokes (RANS) equations.[11,12] Spatial discretization of the governing equations is performed using second-order summation-by-parts operators. Boundary and block interface conditions are enforced weakly through simultaneous approximation terms, which allow C^1 discontinuities in mesh lines at block interfaces. Deep convergence is efficiently achieved using an inexact-Newton phase, while globalization is provided by an approximate-Newton start-up phase. The resulting large, sparse linear system is solved using the flexible generalized minimal residual method with an approximate-Schur parallel preconditioner. The RANS equations are closed using the Spalart-Allmaras one-equation turbulence model. A scalar artificial dissipation scheme[13,14] is used for the cases in this paper, but matrix dissipation[15] can also be used.

C. Gradient Evaluation and Optimization Algorithm

The general optimization problem can be posed as follows:

$$\begin{aligned}
&min && \mathcal{J}(\mathbf{v},\mathbf{q},\mathbf{b}^{(m)}) && \text{(5a)}\\
&w.r.t. && \mathbf{v} && \text{(5b)}\\
&s.t. && \boldsymbol{\mathcal{M}}^{(i)}(\mathbf{A}^{(i)}(\mathbf{v}),\mathbf{b}^{(i)},\mathbf{b}^{(i-1)}) = \boldsymbol{\mathcal{R}}(\mathbf{v},\mathbf{q},\mathbf{b}^{(m)}) = 0, \quad i=1,2,...,m && \text{(5c)}
\end{aligned}$$

where $\mathcal{J}$ is the objective function, $\mathbf{v}$ are the design variables, $\mathbf{b}^{(i)}$ are the volume control-point coordinates at mesh movement increment i, $\boldsymbol{\mathcal{M}}^{(i)}$ are the mesh residuals, and $\mathbf{A}^{(i)}$ are the displaced surface control-point coordinates. The design variables $\mathbf{v}$ are either a subset of $\mathbf{b}^{(m)}$, if using B-spline surface geometry control, or the FFD control points, if using FFD geometry control, and may also include angle of attack. There can be additional linear and nonlinear equality and inequality constraints.

Gradient Evaluation

Gradients are calculated using the discrete-adjoint method at a cost virtually independent of the number of design variables. While it has been shown that gradient-based multistart or hybrid algorithms can be used for multimodal problems,[16] this approach is not taken here. To perform the constrained optimization, the Lagrangian function is introduced:

$$\mathcal{L} \equiv \mathcal{J} + \mathbf{\Lambda}^T \mathbf{c} \tag{6}$$

where $\mathbf{\Lambda}^T = \{\boldsymbol{\lambda}^{(i)}, \psi\}_{i=1}^m$ are the Lagrange multipliers, also called the adjoint variables. For optimality, the Karush-Kuhn-Tucker (KKT) conditions must be satisfied.[17] Once the mesh movement and flow solution have been computed, the resulting flow and mesh adjoint equations must be solved. The flow Jacobian matrix is formed by linearizing its components, including the viscous and inviscid fluxes, the artificial dissipation, the turbulence model, and the boundary conditions. The flow adjoint system is solved using a modified, flexible version of GCROT,[18–20] and the mesh adjoint system is solved using a preconditioned conjugate-gradient method. The final KKT condition gives rise to the objective gradient calculation:

$$\frac{d\mathcal{J}}{d\mathbf{v}} = \frac{\partial\mathcal{J}}{\partial\mathbf{v}} + \sum_{i=1}^{m}(\boldsymbol{\lambda}^{(i)T}\frac{\partial\boldsymbol{\mathcal{M}}^{(i)}}{\partial\mathbf{A}^{(i)}}\frac{\partial\mathbf{A}^{(i)}}{\partial\mathbf{A}^{(m)}}\frac{\partial\mathbf{A}^{(m)}}{\partial\mathbf{v}}) + \psi^T\frac{\partial\boldsymbol{\mathcal{R}}}{\partial\mathbf{v}}. \tag{7}$$

SNOPT

Once the gradients are computed, they are passed to SNOPT (Sparse Nonlinear OPTimizer),[21] a gradient-based optimization algorithm. It can handle both linear and nonlinear constraints, satisfying linear constraints exactly. SNOPT applies a sparse sequential quadratic programming algorithm that approximates the Hessian using a limited-memory quasi-Newton method.

III. Results

A. Case 1: Symmetric Optimization of NACA 0012 Airfoil in Inviscid Transonic Flow

Optimization Problem

The optimization problem is the drag minimization of a modified NACA 0012 airfoil in inviscid, transonic flow. The freestream Mach number is 0.85, and the angle of attack is fixed at 0°, based on Vassberg et al.[22] The design variables are the z-coordinates of the B-spline surface control points. The thickness is constrained to be greater than or equal to the initial airfoil thickness along the entire chord. Since the main challenge of this problem involves minimizing wave drag while satisfying this minimum thickness constraint, nonlinear thickness constraints are applied at specified locations along the airfoil surface, as opposed to the usual "fit-dependent" approach of linearly constraining the surface control points. Consistent with Bisson, Nadarajah, and Dong,[23] the thickness constraints are enforced at 15%, 20%, 22%, 24%, 26%, 29%, and 35% chord, since the optimizer otherwise tries to reduce the airfoil thickness in this region. Satisfaction of the thickness constraint along the rest of the airfoil is verified once the final shape is obtained. Linear symmetry constraints maintain a symmetric airfoil. The problem can be summarized as

$$\begin{aligned} &\text{minimize } C_d \\ &\text{wrt } z \\ &\text{subject to } z \geq z_{\text{baseline}}, \end{aligned}$$

where C_D is the drag coefficient, z is the z-coordinate of a node on the optimized airfoil, and z_{baseline} is the z-coordinate of the corresponding node on the initial geometry.

Initial Geometry

The initial airfoil is a NACA 0012 modified to have a zero-thickness trailing edge. The airfoil is defined by

$$z_{\text{baseline}} = \pm 0.6(0.2969\sqrt{x} - 0.1260x - 0.3516x^2 + 0.2843x^3 - 0.1036x^4),\ x \in [0, 1]. \tag{8}$$

The modification is to the x^4 term coefficient, to allow for a zero-thickness trailing edge.

Grid

A structured H-topology grid around a flat plate of unit chord is inflated using the mesh movement methodology to fit the NACA 0012 section. Two surface patches define the geometry, one on the top and one on the bottom. To establish mesh convergence, four grid levels are considered. Starting with the fine grid, the medium and coarse grids are obtained by removing every second node in the chordwise and normal directions. A superfine grid is obtained by parametric refinement, which doubles the number of nodes in each direction according to a hyperbolic mesh spacing law. The coarse grid is displayed in Figure 1. Since Jetstream was developed for 3D optimization, the airfoil grids are extruded in the spanwise direction. Ten nodes are located along the unit span. Key grid spacing parameters are recorded in Table 1. Flow analysis was performed for the initial geometry fitted with 48 streamwise control points per surface, and the drag coefficient values are recorded in Table 2. Between the fine and superfine grids, the required resolution of 0.1 drag counts is achieved. The coarse, medium, and fine grids are used for optimization.

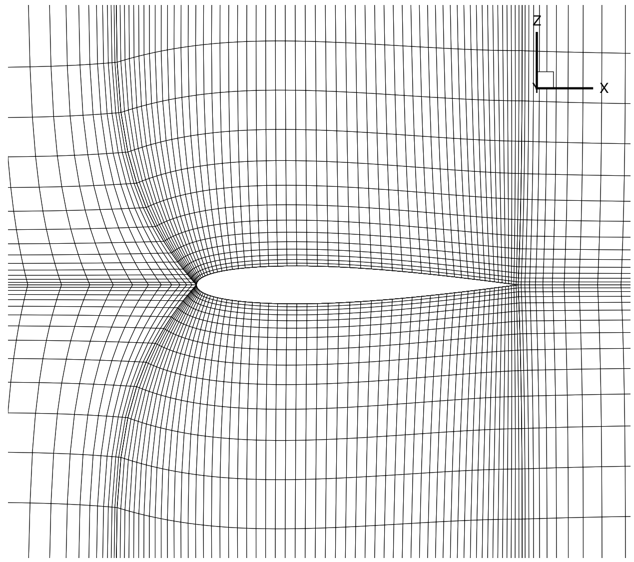

Figure 1: Case 1 - Coarse grid for modified NACA 0012 airfoil

Table 1: Case 1 - Grid parameters for NACA 0012 airfoil grid study

Grid	Nodes (2D)	Off-wall Spacing	Leading-Edge Spacing	Trailing-Edge Spacing
Coarse	12,760	0.008	0.008	0.008
Medium	49,020	0.004	0.004	0.004
Fine	192,100	0.002	0.002	0.002
Superfine	768,400	0.001	0.001	0.001

Table 2: Case 1 - Results of grid study for initial NACA 0012 airfoil

Grid Level	Nodes (2D)	C_D (Counts)
Coarse	12,760	461.299
Medium	49,020	457.598
Fine	192,100	457.327
Superfine	768,400	457.327

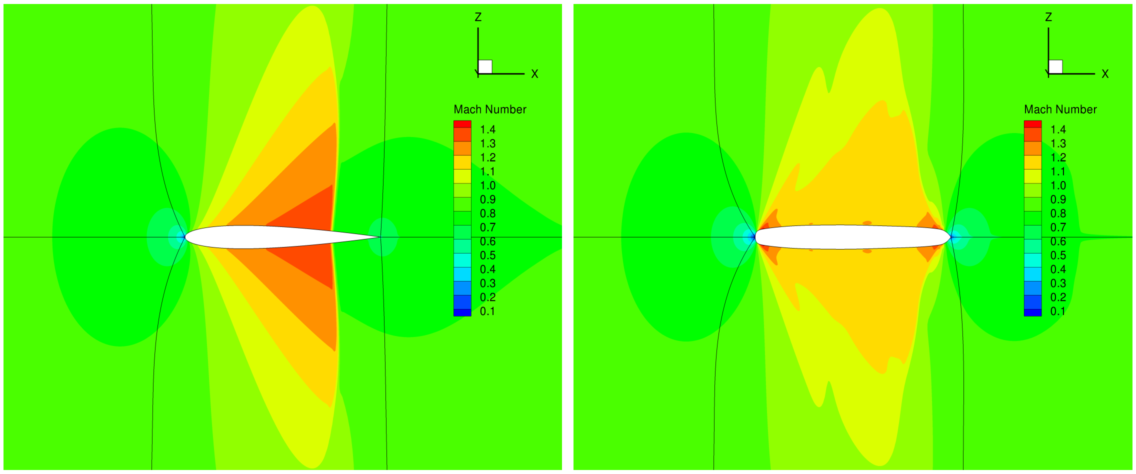

(a) Initial

(b) Final

Figure 2: Case 1 - Comparison of Mach contours for initial NACA 0012 airfoil and final optimized shape with 9 design variables

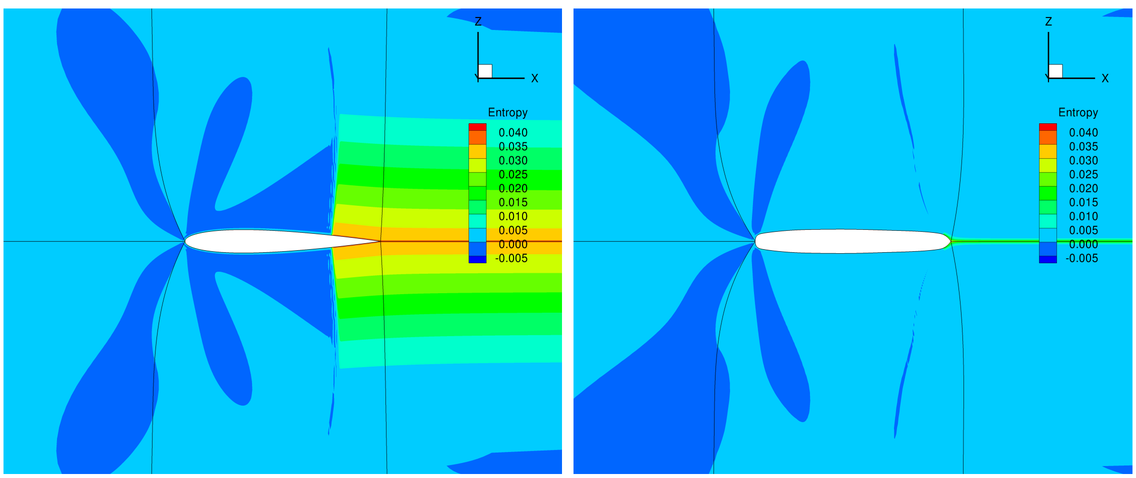

(a) Initial

(b) Final

Figure 3: Case 1 - Comparison of entropy contours for initial NACA 0012 airfoil and final optimized shape with 9 design variables

Optimization Results

B-spline surface geometry control was used for this case. Mach and entropy contours are displayed in Figures 2 and 3, respectively, for the initial and optimized geometries using 9 design variables per surface on the fine mesh. Strong shocks extending far into the flow field are evident on the initial geometry. Due to the thickness constraint, the optimizer thickens the airfoil, creating a relatively flat surface that delays the pressure recovery. Weaker shocks, which do not extend as far into the flow field, occur near the trailing edge of the optimized airfoil. The optimized geometry is quite blunt, since the optimizer is exploiting the fact that the Euler equations cannot correctly model the physics of flow separation, but it is worth noting that RANS analysis would likely show significant separation. The design problem is meant to be more of a challenging academic problem than a practical one.

Optimizations were conducted to investigate the effect of design space dimensionality. The number of B-spline control points used to parameterize each surface ranged from 5 to 13, but since the leading and trailing edge control points were fixed, the corresponding number of design variables for each surface ranged from 3 to 11. Figure 4 plots airfoil surfaces and corresponding pressure coefficients for the initial and final

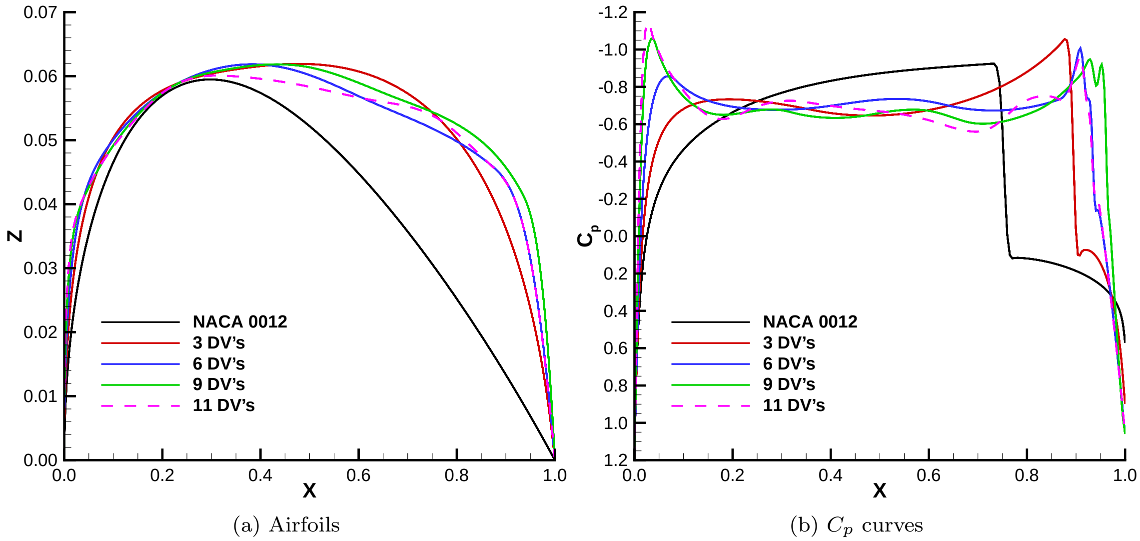

Figure 4: Case 1 - Comparison of initial and final airfoil shapes and corresponding pressure distributions, optimized and analyzed on fine mesh level

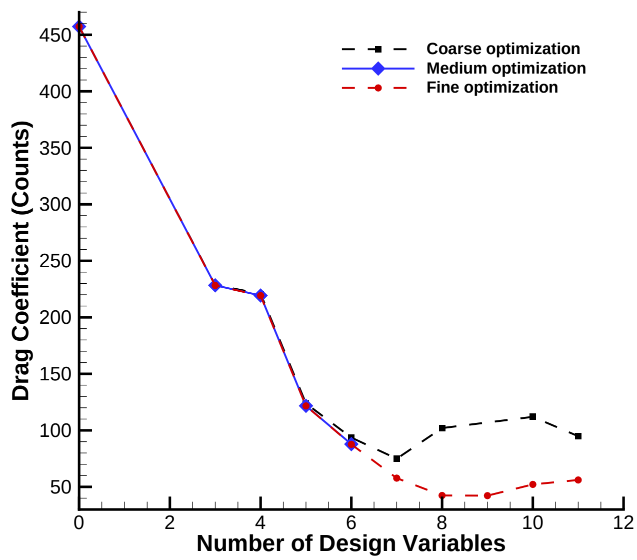

Figure 5: Case 1 - Drag comparison of final geometries from optimizations on coarse, medium, and fine grid levels, evaluated on the fine mesh

geometries from optimizations conducted on the fine mesh, with 3, 6, 9, and 11 design variables on each surface. With more design variables, the optimizer has more freedom to create blunter leading and trailing edges. The suction peak becomes more abrupt; the shock is weakened and pushed further downstream.

To investigate the effect of grid density, optimizations were conducted using the coarse, medium, and fine grids. The drag coefficient of the final geometries evaluated on the fine mesh is plotted against the number of design variables in Figure 5. The drag from the medium- and fine-mesh optimizations is nearly identical. The drag from the fine-mesh optimizations is consistently lower than that obtained from the coarse-mesh optimizations. The difference is below 2 counts for 3 to 5 design variables, but is approximately 60 counts for 8 design variables. The lowest drag obtained is from the fine-mesh optimization using 9 design variables per surface: 42.24 drag counts, a 91% reduction from the initial geometry.

Figure 6 shows the convergence history for 3, 6, 9, and 11 design variables. Optimality, which is a measure of convergence of the optimization to a local minimum, has a convergence tolerance set at $1{\times}10^{-7}$. Feasibility, which is a measure of how well the nonlinear constraints are satisfied, has a convergence tolerance set at $1{\times}10^{-6}$. The optimizer has no difficulty satisfying the nonlinear thickness constraints and therefore

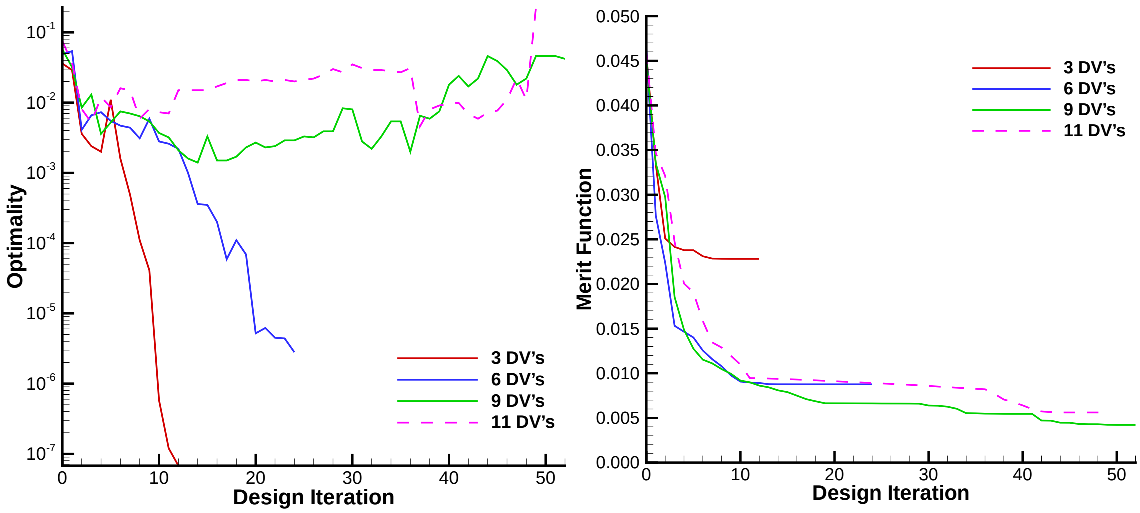

(a) Optimality

(b) Merit function

Figure 6: Case 1 - Typical convergence histories for NACA 0012 airfoil optimization

Table 3: Case 1 - Grid study for NACA 0012 airfoil optimizations on coarse grid

	Drag Coefficient, C_d (Counts)		
Design Variables	**Coarse Flowsolve**	**Medium Flowsolve**	**Fine Flowsolve**
3	221.34	224.99	228.61
4	217.31	218.72	220.83
5	145.93	124.05	123.8
6	138.87	96.53	93.61
7	143.23	82.48	74.93
8	151.87	108.21	102.03
9	149.98	116.74	failed
10	143.64	113.41	111.96
11	136.17	97.46	94.88

feasibility is not plotted. As the number of design variables increases, the optimizer has a more difficult time reducing optimality. Drag, however, is still reduced over the course of the optimizations, as shown by the merit function plots. The higher drag for cases with more design variables, such as the 11 design variable case shown, is likely associated with poor optimizer convergence.

Tables 3 to 5 display the final drag values for optimizations conducted on the coarse, medium, and fine grid levels, respectively. Each final geometry was analyzed on all three grid levels. While the grid study on the initial geometry gave a drag difference of less than a count between the medium and fine grids, the differences for the final geometries are considerably larger. Several failed flow solves occur during fine-mesh analysis, suggesting that the fine grid is resolving difficult flow features not observed with the coarser spacing, making the problem more difficult to converge.

More concerning, however, is the presence of converged solutions with finite lift, indicative of the presence of non-unique solutions. Note that all of the geometries are symmetric. The presence of non-unique solutions may be due to the fact that the Euler equations are ill-suited for such bluff bodies. As previously stated, RANS analysis would likely give considerable flow separation, and unsteady analysis would perhaps give unsteady flow features.[24] Hence, forcing the flow to remain tangent to the highly "blunt" trailing edge is unphysical, and this may be contributing to the ill-posedness of the problem.

To examine the lifting solutions further, the pressure distribution from the fine-mesh analysis of the final geometry from the 10 design variable case on the medium mesh is displayed in Figure 7. A double shock is observed on the lower surface and a single shock on the upper surface, consistent with the double shocks

Table 4: Case 1 - Grid study for NACA 0012 airfoil optimizations on medium grid

	Drag Coefficient, C_d (Counts)		
Design Variables	**Coarse Flowsolve**	**Medium Flowsolve**	**Fine Flowsolve**
3	221.09	224.72	228.33
4	216.61	217.28	219.3
5	146.88	122	121.71
6	152.55	91.71	87.82
7	168.94	68.36	nonzero lift
8	nonzero lift	60.67	nonzero lift
9	200.39	68.25	failed
10	nonzero lift	69.64	nonzero lift
11	nonzero lift	67.56	failed

Table 5: Case 1 - Grid study for NACA 0012 airfoil optimizations on fine grid

	Drag Coefficient, C_d (Counts)		
Design Variables	**Coarse Flowsolve**	**Medium Flowsolve**	**Fine Flowsolve**
3	220.96	224.58	228.2
4	216.68	217.23	219.22
5	146.42	121.96	121.62
6	143.77	91.86	87.63
7	169.32	68.86	57.74
8	nonzero lift	58.99	42.39
9	229.38	nonzero lift	42.24
10	219.25	77.93	52.18
11	244.86	87.14	56.1

observed by Jameson et al.,[25] who analyzed similar bluff airfoils under similar flow conditions. Mach contours and streamlines are displayed in Figure 8, with a zoomed view of the trailing edge region. Recirculation is observed. Examination of other lifting solutions shows similar flow features.

Flow analysis was conducted on the fine mesh for the geometry obtained from fine-mesh optimization with 9 design variables, sweeping down from an initial Mach number of 0.8524 to 0.8493, and then back up again. Each flow solution was initialized using the previous converged solution. Figure 9 shows that hysteresis behaviour is observed; non-unique solutions occur near the operating condition of Mach 0.85.

A closer look at the optimization histories shows that non-unique, lifting solutions of the same nature were produced during the optimizations. For example, C_l and C_d are plotted during the 10 design variable optimization on the medium mesh in Figure 10. Lifting solutions give high drag values that the optimizer does not expect from the gradient information. Such instances are detrimental to the optimization convergence. In addition, more design variables means more geometric flexibility, but with this comes an increased ability to find non-unique solutions as the airfoil becomes increasingly bluff. So while general intuition and experience says that more design variables means increased ability to reduce drag, this does not necessarily hold true for this ill-posed problem.

Finally, the convergence histories of converged and failed flow solutions are examined. Samples of lift coefficient and residual histories for a zero-lift converged solve, a non-zero-lift converged solve, and a failed solve are shown in Figure 11. These examples are typical of other examined cases. Throughout the converged zero-lift case, the lift remains close to zero. It seems that the problem is more difficult to fully converge if a lifting solution appears. The number of converged lifting and failed flow solves during an optimization generally increases as the number of design variables is increased, resulting in poor optimization convergence.

To sum up the results for this case, the airfoil with lowest drag is obtained by optimizing on the fine mesh with 9 design variables on each surface. A C_d of 42.2 drag counts is computed on the fine mesh.

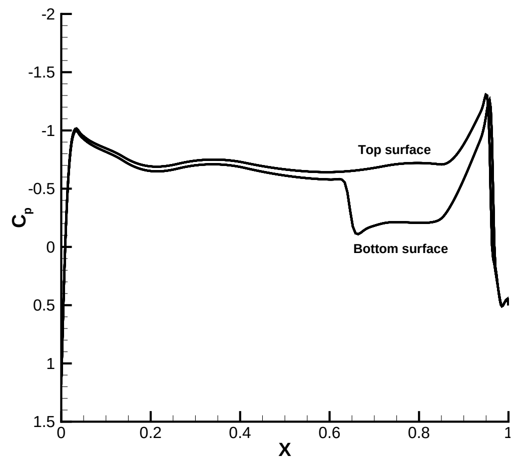

Figure 7: Case 1 - Pressure coefficient for final geometry from 10 design variable optimization on medium mesh, analyzed on fine mesh

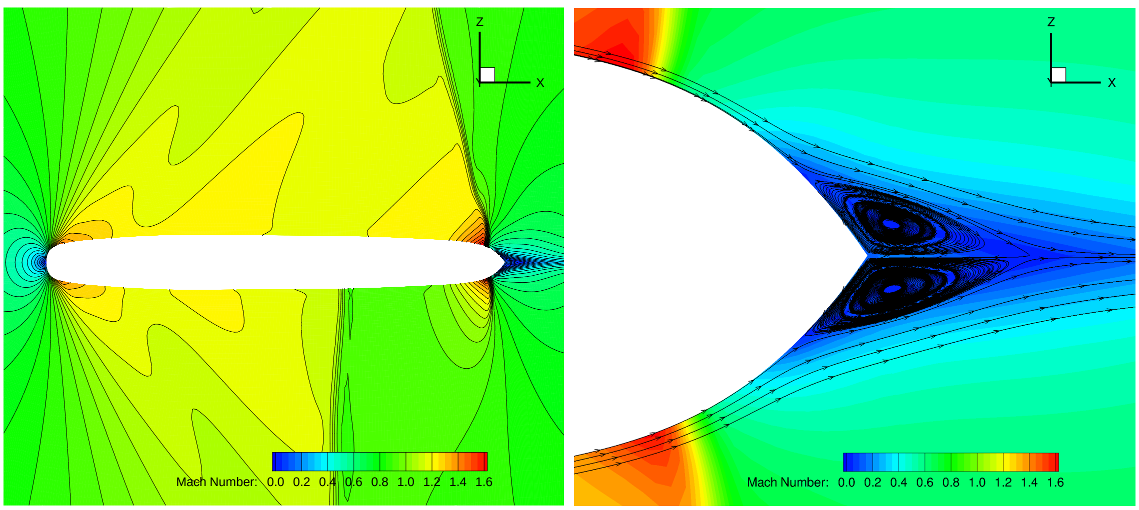

(a) Mach contours about airfoil

(b) Streamlines in trailing edge region

Figure 8: Case 1 - Mach contours and streamlines for final geometry from 10 design variable optimization on medium mesh, analyzed on fine mesh

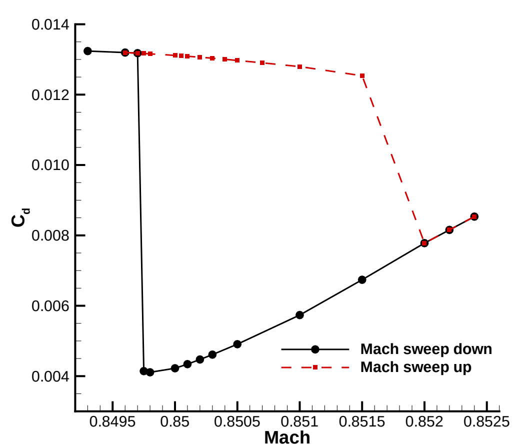

Figure 9: Case 1 - Drag coefficient hysteresis over Mach number

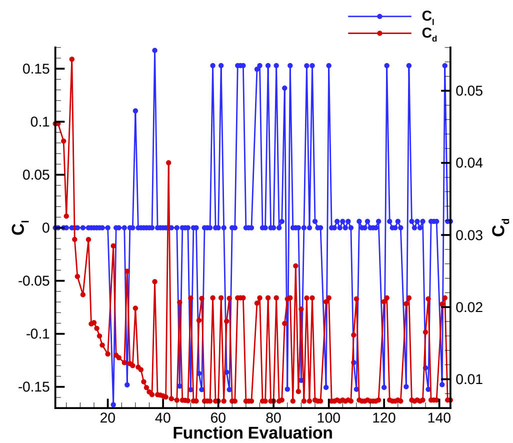

Figure 10: Case 1 - Functional history during 10 design variable optimization on medium mesh

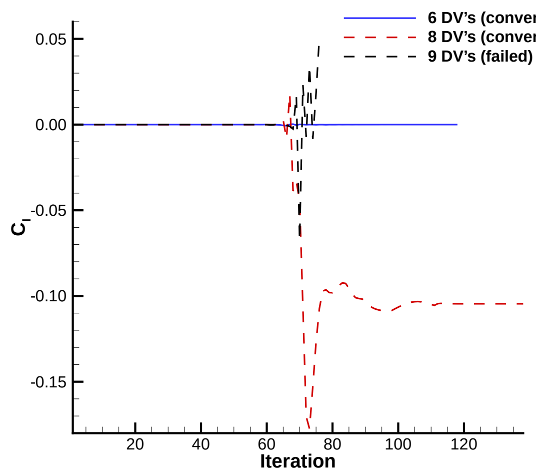

(a) Lift coefficient

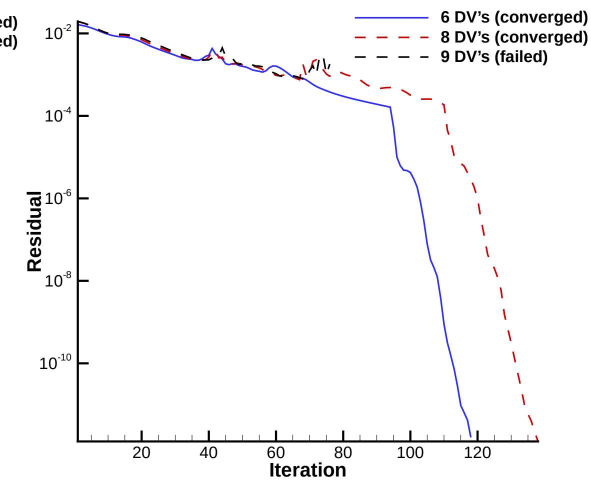

(b) Residual norm

Figure 11: Case 1 - C_l and residual histories for fine-mesh flow solves on geometries from optimizations on medium mesh using 6, 8, and 9 design variables per surface

B. Case 2: Optimization of RAE 2822 Airfoil in Viscous Transonic Flow

Optimization Problem

The optimization problem is the drag minimization of the RAE 2822 airfoil in viscous, transonic flow. The freestream Mach number is 0.734, and the Reynolds number is 6.5 million. The design variables are the z-coordinates of the B-spline surface or FFD control points, as well as the angle of attack. The lift coefficient is constrained to 0.824 and the moment coefficient about the quarter-chord must be no less than -0.092. The minimum airfoil area is the initial airfoil area. Though not required by the test description, a minimum thickness constraint of 25% of the initial thickness is enforced to maintain a realistic design and prevent control point cross-over. The problem can be summarized as

$$\begin{aligned} \text{minimize } & C_d \\ \text{wrt } & z,\ \alpha \\ \text{subject to } & C_l = 0.824 \\ & C_m \geq -0.092 \\ & A \geq A_{\text{baseline}} \end{aligned}$$

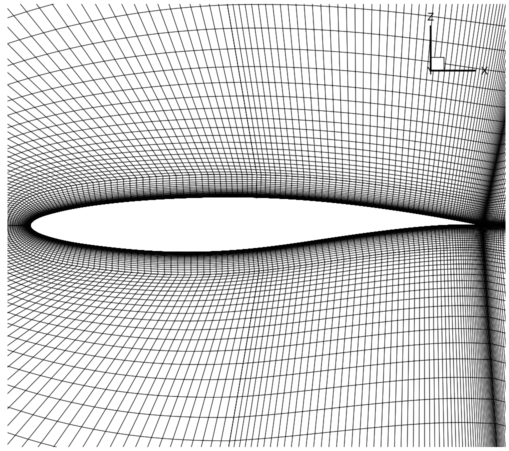

Figure 12: Case 2 - Coarse grid of RAE 2822 used for optimizations

Table 6: Case 2 - Grid parameters for RAE 2822 airfoil grid study

Grid	Nodes (2D)	Off-wall Spacing	Leading-Edge Spacing	Trailing-Edge Spacing
Coarse	47,824	3.7×10^{-6}	1×10^{-3}	1×10^{-3}
Medium	187,792	1.8×10^{-6}	5×10^{-4}	5×10^{-4}
Fine	748,064	8.7×10^{-7}	2.5×10^{-4}	2.5×10^{-4}
Superfine	3,016,832	4.3×10^{-7}	1.25×10^{-4}	1.25×10^{-4}
Finest	12,067,328	2.1×10^{-7}	6.25×10^{-5}	6.25×10^{-5}

where C_d, C_l, and C_m are the drag, lift, and moment coefficients, respectively, and A and A_{baseline} are the optimized and initial airfoil areas, respectively.

Initial Geometry

The initial geometry is the RAE 2822 airfoil. The coordinates are obtained from the UIUC Airfoil Coordinates Database.[b]

Grid

A C-topology grid is used for this case, and the optimization grid is displayed in Figure 12. To establish grid convergence, the optimization grid was repeatedly refined by a factor of two in each direction, giving the grid family with parameters shown in Table 6. The locations of the new nodes added during refinement were determined according to a consistent hyperbolic mesh spacing law, giving a consistent grid family.

On the optimization grid, an angle of attack of 3.119° satisfies the desired C_l of 0.824. Coefficients of drag, lift, and moment, as well as average y^+, are reported in Table 7. The desired drag resolution of 0.1 counts is achieved between the two finest grid levels, but lift is not within 0.1×10^{-4}.

To evaluate the accuracy of the finest grid level, numerical results were obtained to compare to the commonly referenced experimental results for the RAE 2822 - Case 9.[26] The flow conditions for the analysis were set to match the experimental flow conditions: a Mach number of 0.73, Reynolds number of 6.5 million, and corrected wind tunnel angle of attack of 2.79°. The experimental results give a normal force coefficient C_n of 0.803, a C_d of 0.0168, and a C_m of -0.099. The analysis on the finest grid level gives a C_n of 0.802, a C_d of 0.0167, a C_m of -0.091. In addition, pressure coefficients are compared in Figure 13. The computed results agree well with experiment.

[b] `http://m-selig.ae.illinois.edu/ads/coord_database.html`

Table 7: Case 2 - Results of grid study for initial RAE 2822 airfoil

Grid Level	y^+	C_l	C_d (Counts)	C_m
Coarse	0.71	0.8240	234.44	-0.0908
Medium	0.32	0.8439	228.41	-0.0932
Fine	0.15	0.8501	229.14	-0.0944
Superfine	0.076	0.8519	229.61	-0.0947
Finest	0.038	0.8524	229.69	-0.0948

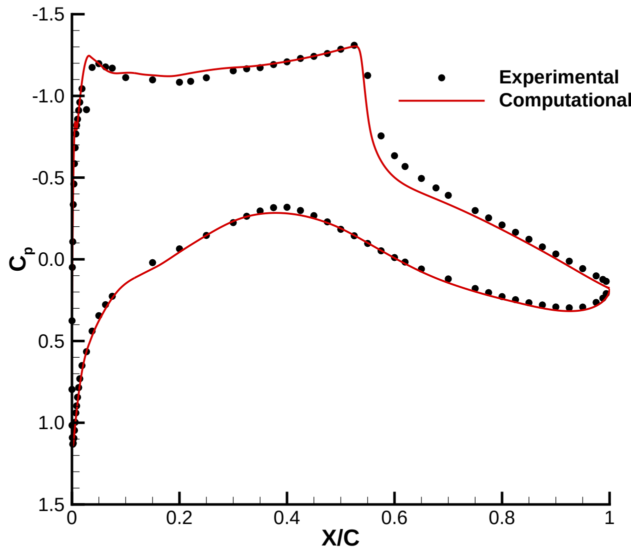

Figure 13: Case 2 - Comparison of experimental and computed pressure coefficient on finest grid level

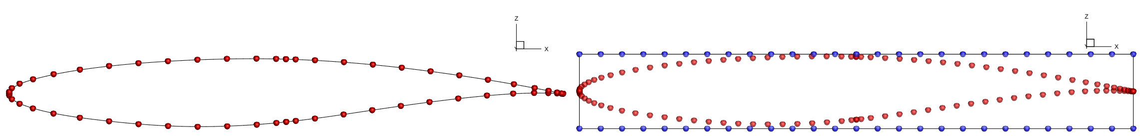

(a) Initial B-spline optimization setup

(b) Initial FFD optimization setup

Figure 14: Case 2 - Initial 27 chordwise geometric design variables per surface for B-spline surface (red points) and FFD (blue points) control

Optimization Results

To examine the effect of design space dimensionality, optimizations with B-spline surface control were conducted with from 7 to 37 chordwise design variables on the top and bottom surfaces, and optimizations with FFD control were conducted with from 5 to 37 design variables on the top and bottom of the FFD volume. For the FFD cases, 53 chordwise surface control points parameterize the top and bottom surfaces of the airfoil. The initial design variables for both 27 design variable cases are displayed in Figure 14. Angle of attack is also a design variable.

Lift, drag, and moment coefficient values, along with angle of attack, for the final geometries evaluated on the optimization grid are plotted in Figure 15. Although monotonic drag reduction is observed when the number of design variables is increased for a relatively low number of design variables, the drag reduction performance degrades for higher numbers of design variables. Noticeable differences in the final geometries are also evident, as shown in Figure 16. The 17 B-spline surface design variable case gave the lowest drag among the B-spline control optimizations, and the 11 FFD design variable case gave the lowest drag among the FFD control optimizations. The 27 design variable cases did not perform well. The superior geometries exhibit

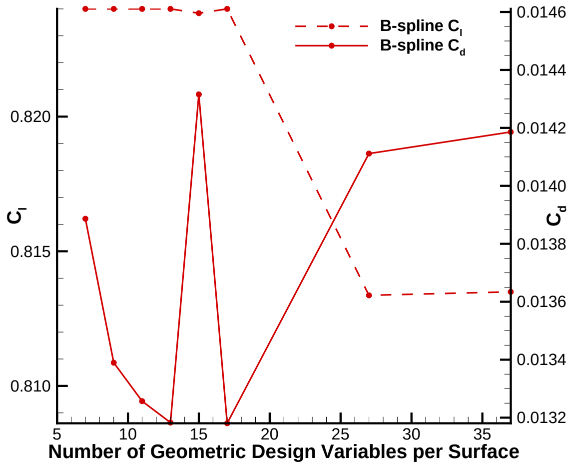

(a) B-spline control C_l and C_d results

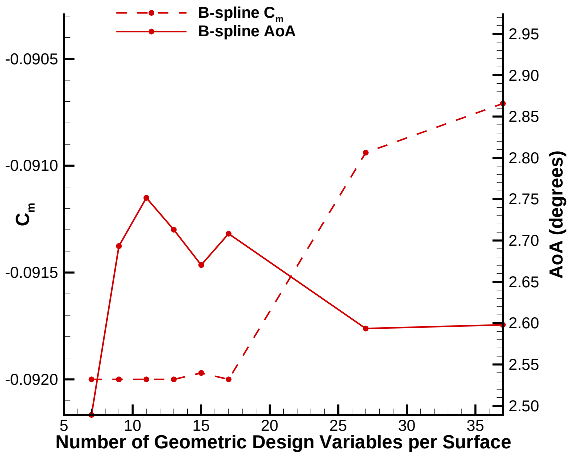

(b) B-spline control C_m and α results

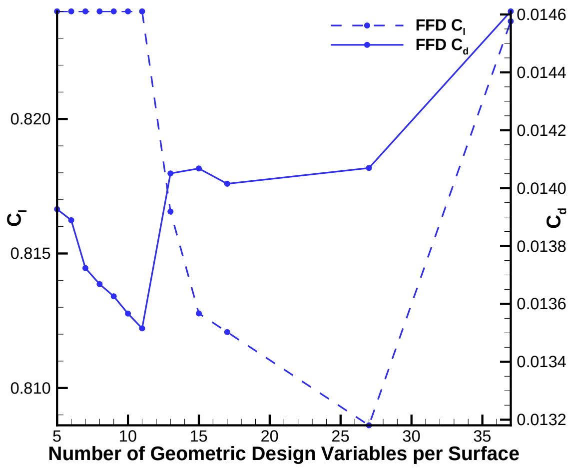

(c) FFD control C_l and C_d results

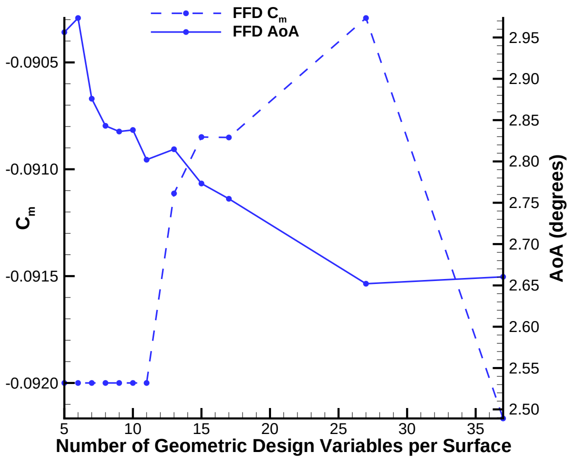

(d) FFD control C_m and α results

Figure 15: Case 2 - Coarse-mesh optimization functionals and angle of attack

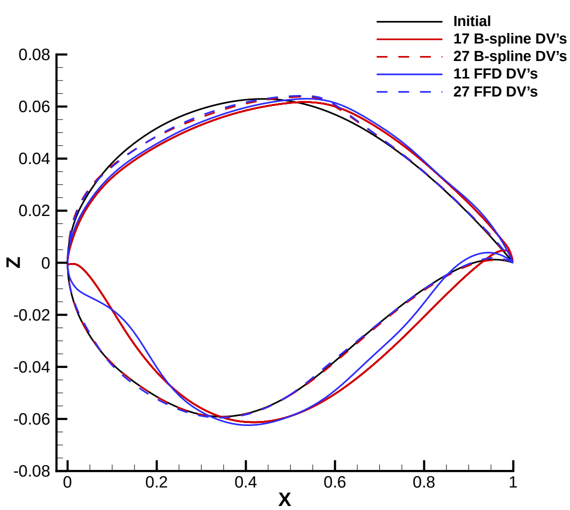

(a) Airfoil shapes

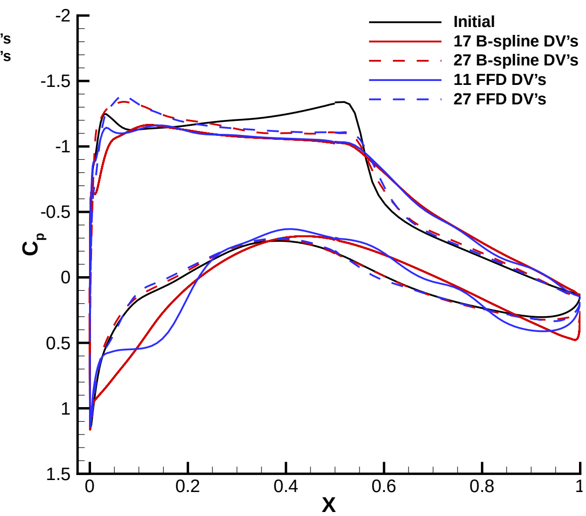

(b) C_p curves

Figure 16: Case 2 - Airfoil shapes and pressure coefficients

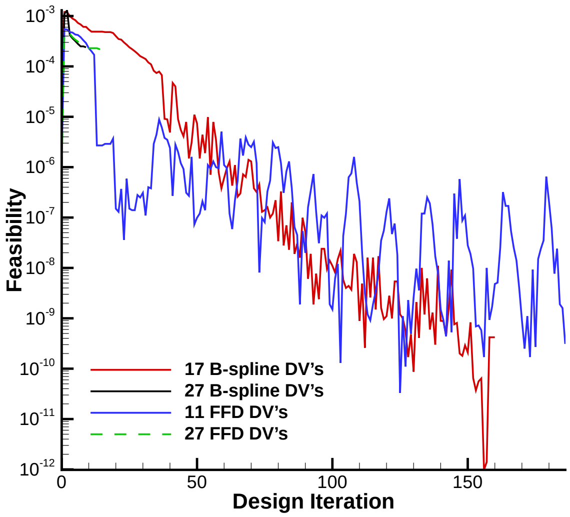

(a) Feasibility

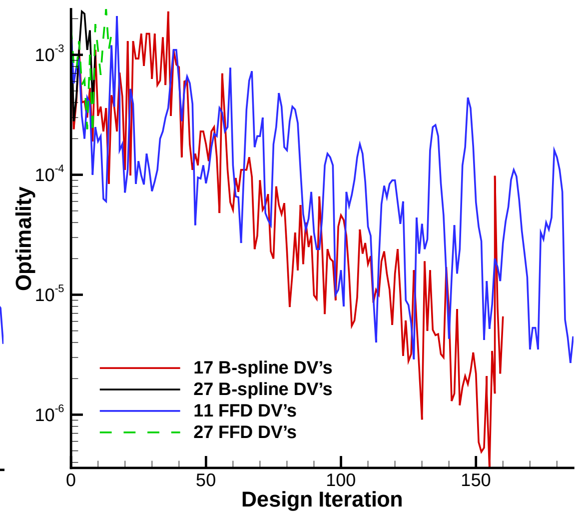

(b) Optimality

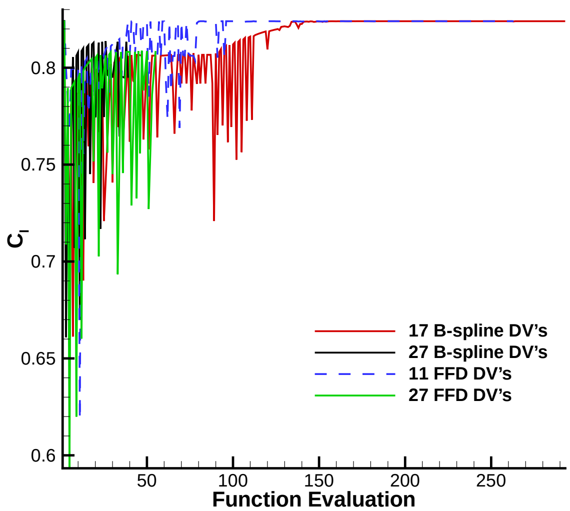

(c) C_l

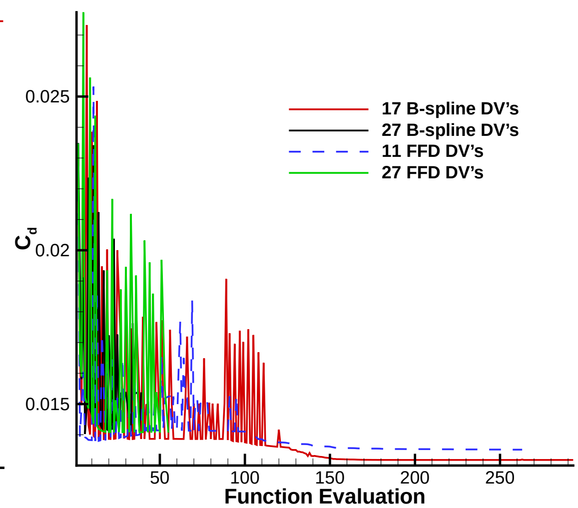

(d) C_d

Figure 17: Case 2 - Convergence histories

small leading-edge radii and highly cambered trailing edges. Due to the uniform chordwise distribution of FFD control points used, less control is offered at the leading and trailing edges, and the geometric features are not as pronounced.

Figure 17 compares the optimization convergence histories of the B-spline surface control cases with 17 and 27 design variables and FFD control cases with 11 and 27 design variables. Figures 17a and 17b plot feasibilility and optimality as a function of design iterations. The optimality and feasibility tolerances are set at 1×10^{-7} and 1×10^{-6}, respectively. The 17 B-spline and 11 FFD design variable cases show far superior optimization covergence. The 27 design variable cases ran very few iterations and do not satisfy the C_l constraint. Figures 17c and 17d display lift and drag coefficient histories versus function evaluations. While the 17 B-spline and 11 FFD design variable cases exhibit drag spikes early on in the optimization, these spikes disappear and monotonic drag reduction is observed. In contrast, the 27 design variable cases continue to periodically produce designs with dramatic drag increases and lift decreases.

The geometry at iteration 36 from the 27 surface design variable case was re-analyzed using three different convergence paths to steady state. Two distinct fully converged solutions were obtained, giving (C_l, C_d) pairs of (0.7830, 0.01632) and (0.8133, 0.01411). Unsteady RANS analysis demonstrated that the two steady solutions are physically stable. The pressure distributions are displayed in Figure 18. Note that the single-block solution shown will be discussed later. Interestingly, the discrepancies in the solutions are quite

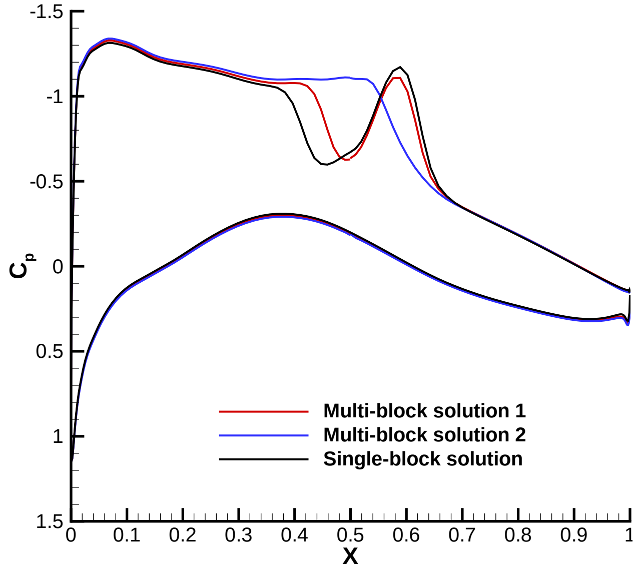

Figure 18: Case 2 - Pressure distributions for non-unique solutions

localized. The double pressure recovery features are consistent with curves observed for this case by LeDoux et al.[27]

The presence of non-unique solutions in the design space explains the poor convergence of some of the optimizations. Analyses of other geometries using the different convergence paths were conducted, and while non-unique solutions were not always observed, they were found to be common. Gradient-based optimization does not work well with design spaces that are not smooth, since search directions provided by gradient calculations can lead to unexpected spikes in the design space. For reasons not fully understood, the combination of geometries produced by the optimizations with these flow conditions yield ill-posed problems. It also appears that the occurence of non-unique solutions is a greater problem with more design variables, since the increased geometric flexibility gives more freedom to fall into these undesired regions in the design space. While one would expect that the cases with more design variables should have the geometric flexibility to produce the geometries resembling the fewer design variable geometries, it seems that with more design variables and occurences of non-unique solutions, the optimizer is more prone to get stuck.

The grid for geometry number 36 from the 27 surface design variable case was converted from a multi-block grid to single-block grid. Analysis of the single-block grid in Jetstream with different convergence paths all gave a single solution, with lift and drag coefficient values of 0.7570 and 0.01711, respectively. In addition, this single solution is distinct from both multi-block solutions, as shown in the pressure coefficient plots of Figure 18. A double pressure recovery is observed. Several other geometries that gave non-unique solutions with multi-block meshes were analyzed with a single-block mesh, and all gave unique solutions. This suggests that the non-uniqueness may be triggered by the treatment of block interfaces; however, further investigation of this hypothesis has yet to be conducted.

While the occurence of non-unique solutions possibly arising from the multi-block treatment of the flow solver is a concern, it is worth mentioning that Jetstream and its 2D predecessor Optima2D[28] have been been used to solve many transonic optimization problems in the past, without observing non-uniqueness or the associated convergence difficulties.[12,29–33] The phenomenon only seems to arise under very specific flow conditions that produce a sufficiently ill-posed problem. When the flow conditions are adjusted to a Mach number of 0.75 and lift coefficient constraint of 0.6, non-unique solutions are no longer observed.

To sum up the results for this case, the airfoil with lowest drag is obtained with 17 design variables on each surface. To establish grid convergence for this final geometry, a grid refinement study is performed at the final angle of attack of 2.708°, and the results are recorded in Table 8. A C_d of 119.22 drag counts is computed on the finest mesh.

Table 8: Case 2 - Results of grid study for optimized RAE 2822 airfoil

Grid Level	y^+	C_l	C_d (Counts)	C_m
Coarse	0.75	0.8240	131.81	-0.0920
Medium	0.34	0.8415	121.88	-0.0945
Fine	0.17	0.8436	120.11	-0.0948
Superfine	0.081	0.8433	119.54	-0.0946
Finest	0.041	0.8423	119.22	-0.0944

C. Case 3: Twist Optimization of a Rectangular Wing in Inviscid Subsonic Flow

Optimization Problem

The optimization problem is the drag minimization of a rectangular wing with zero-thickness trailing edge NACA 0012 sections in inviscid, subsonic flow. The freestream Mach number is 0.5. The design variables are the twist of sections along the span about the trailing edge. Twist is performed by allowing the z-coordinates of the B-spline surface or FFD control points to vary under linear constraints, thus linearly shearing the sections. The twist at the root section is allowed to vary, while the angle of attack is fixed. The target lift coefficient is 0.375. The twist distribution should produce a lift distribution close to elliptical and an efficiency factor close to unity. The problem can be summarized as

$$\begin{aligned} &\text{minimize } C_D \\ &\text{wrt } \gamma(y) \\ &\text{subject to } C_L = 0.375 \end{aligned}$$

where C_D and C_L are the drag and lift coefficients, respectively, and $\gamma(y)$ is the twist distribution along the span.

Initial Geometry

The initial geometry is a rectangular, planar wing with NACA 0012 sections. The trailing edge is sharp. The semi-span is $3.06c$, with the last 0.06 leading to a pinched tip. Although the tip geometry is not quite consistent with the case description, which specifies a rounded tip, the main purpose of optimizing the twist distribution is still maintained.

Grid

An H-topology grid was used for this case. It is a flat plate grid that is inflated to the NACA 0012 section using the mesh movement algorithm. The optimization level mesh on the aerodynamic surface and symmetry plane is displayed in Figure 19. To establish grid convergence, the optimization level grid is refined by a factor of 2, 3, and 4 in each direction, giving the grid family with parameters shown in Table 9. The locations of the new nodes added during refinement were determined according to the hyperbolic mesh spacing law.

On the optimization grid, an angle of attack of 4.2040° gives the desired C_L of 0.375 and is used for the grid study. Coefficients of drag and lift, as well as span efficiency factor, are plotted in Figure 20 for the different grid levels. The medium and superfine grids, which differ in grid size by a factor of 2 in each direction, give C_D values within 1 drag count of each other. The three finest grid levels appear to be in the asymptotic region, as the behaviour of both functionals is linear with respect to $N^{-2/3}$, where N is the total number of grid nodes, and the spatial discretization is second-order. The span efficiency factor of the initial geometry is very close to unity on the coarse mesh, but the refinement study shows there is in fact room for improvement. The superfine mesh is chosen for refined analysis of the optimized geometries.

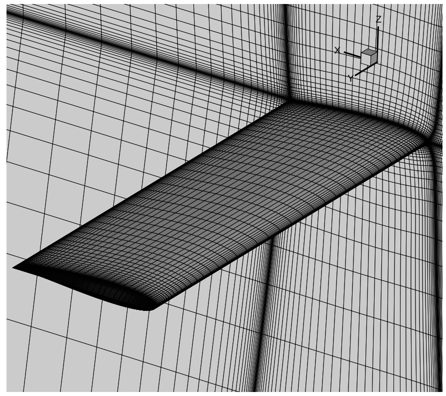

Figure 19: Case 3 - Mesh for aerodynamic surface and symmetry plane

Table 9: Case 3 - Grid parameters for the rectangular planar wing grid study

Grid	Nodes	Off-wall Spacing	Leading-Edge Spacing	Trailing-Edge Spacing
Coarse	1,361,976	3×10^{-3}	3×10^{-3}	3×10^{-3}
Medium	10,895,808	1.5×10^{-3}	1.5×10^{-3}	1.5×10^{-3}
Fine	36,773,352	1×10^{-3}	1×10^{-3}	1×10^{-3}
Superfine	87,166,464	7.5×10^{-4}	7.5×10^{-4}	7.5×10^{-4}

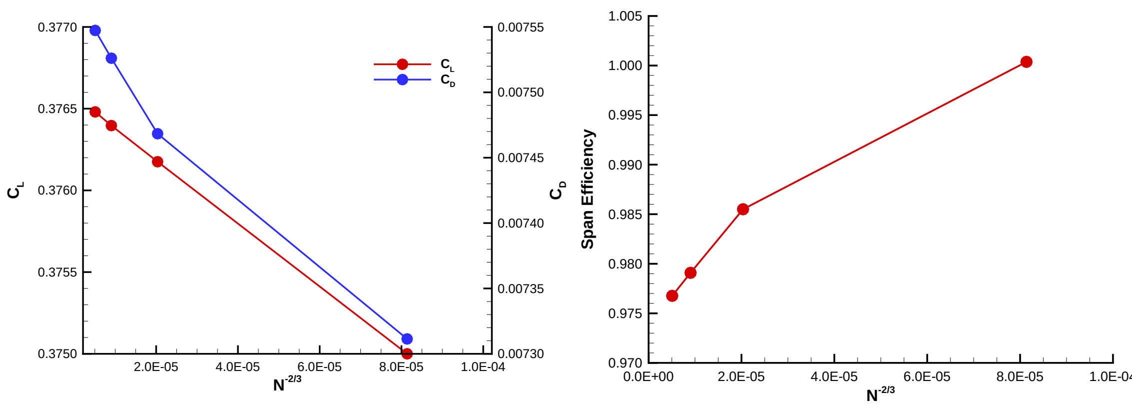

(a) Lift and drag coefficients

(b) Span efficiency

Figure 20: Case 3 - Lift coefficient, drag coefficient, and span efficiency factor evaluated on the different grid levels for the initial geometry

Optimization Results

The optimizations were conducted on the coarse mesh with different numbers of B-spline surface and FFD design variables. For the B-spline surface optimizations, the twist for the spanwise stations on the tip patches was constrained to be a linear extrapolation of the twist between the two adjacent stations on the inboard patches. This was to prevent the optimizer from exploiting the surface control point clustering at the tip to create a non-planar feature. This was not necessary for the FFD optimizations since the FFD spanwise stations were sufficiently (uniformly) spaced out along the span. All of the optimizations were successful, reaching feasibility and optimality tolerances of 1×10^{-6} and 1×10^{-7}, respectively. For example, the feasibility, optimality, and merit function histories for the optimization with 10 B-spline surface design variables are displayed in Figure 21. The final geometries are re-analyzed on the superfine mesh, with the angle of attack adjusted in each case to satisfy the C_L constraint of 0.375. The drag coefficients and span

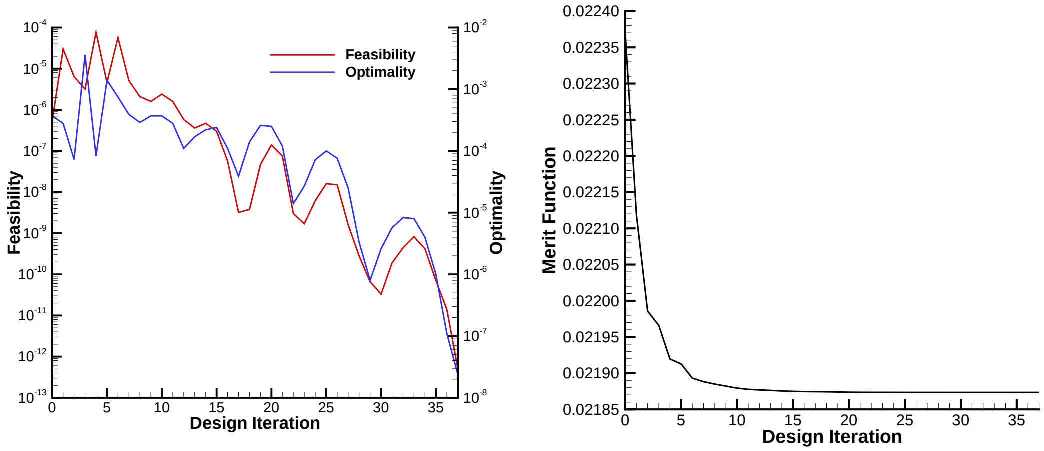

(a) Feasibility and optimality histories

(b) Merit function history

Figure 21: Case 3 - SNOPT convergence history for the optimization with 10 B-spline surface design variables

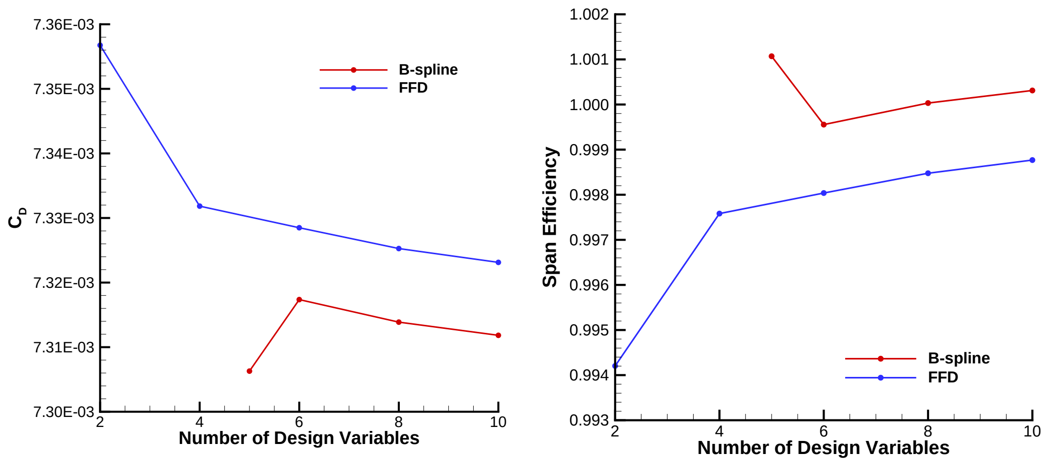

(a) Drag coefficient

(b) Span efficiency

Figure 22: Case 3 - Optimized drag coefficient and span efficiency evaluated on the superfine mesh for different numbers of design variables

efficiency factors are plotted in Figure 22. All of the optimized span efficiencies are very close to unity, and the two efficiencies that exceed unity are a reflection of the fact that linear aerodynamic theory does not consider all the effects of the nonlinear Euler equations. As expected, the spanwise lift distributions are close to elliptical. For example, the lift distributions obtained with 10 design variables are compared to the initial and elliptical distributions in Figure 23. The only noticeable deviation from elliptical occurs at the tip and is attributed to mesh effects at the tip.

D. Case 4: Twist and Section Optimization of CRM Wing in Turbulent Transonic Flow

Optimization Problem

The problem is the drag minimization of the wing geometry extracted from the Common Research Model (CRM) wing-body configuration from the Fifth Drag Prediction Workshop.[4] The goal is to optimize the sectional shape and twist to minimize drag at a lift coefficient of $C_L = 0.5$ at a Mach number of 0.85 and a Reynolds number of 5 million. The design variables are the z-coordinates of either the B-spline surface control points or the FFD volume control points, in addition to the angle of attack. The B-spline points on the trailing edge of the wing are fixed to permit arbitrary twist, except for the root where the leading edge control point is also fixed. The optimization problem is specified as

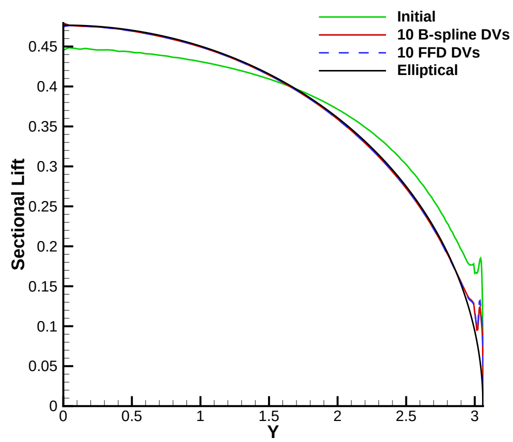

Figure 23: Case 3 - Lift distributions from the 10 design variable optimizations, analyzed on the superfine mesh, compared to initial and elliptical distributions

$$\begin{aligned} \text{minimize } & C_D \\ \text{wrt } & z, \alpha \\ \text{subject to } & C_L = 0.500 \\ & C_M \geq -0.17 \\ & V \geq V_{\text{baseline}} \\ & z \geq 0.25 \times z_{\text{baseline}}. \end{aligned}$$

Initial Geometry

The wing geometry is scaled by the mean aerodynamic chord of 275.8 inches and translated such that the origin is at the root leading edge. Moments are calculated about the point (1.2077, 0, 0.007669). All aerodynamic force coefficients are calculated using a reference area of $S_{\text{ref}} = 3.407$ squared reference units. The initial volume, V_{baseline}, is 0.2617 cubed reference units. The wing surface is divided into three spanwise sections, with each section consisting of two chordwise patches. An additional two patches are used for the leading edge and blunt trailing edge. Finally, the wing tip cap is modelled by two patches, giving a total of 20 surface patches. For B-spline surface geometry control, the leading-edge patches are parameterized by 5 points in the streamwise and spanwise directions, while all other patches have 9 points in the streamwise direction and 5 points in the spanwise direction. This gives a total of 15 spanwise design sections, each controlled by 35 points. For FFD geometry control, all patches are parameterized by 17 by 17 control points and are embedded inside two FFD volumes joined at the wing crank. The FFD volumes give 15 spanwise design sections, each controlled by 17 chordwise FFD control points on the top and bottom, giving a similar number of design variables to the B-spline surface control setup. The FFD points are clustered towards the leading and trailing edges according to a cosine distribution.

Grid

The grid is generated in ICEM CFD and uses an O-O mesh topology. Table 10 shows the information on the different grid levels. The grid is refined in all directions by factors of 2 and 4 to give three levels in total. Figure 24 shows the surface and symmetry planes for the the computational mesh as well as the optimization B-spline surface. Figure 25 shows the results of the grid convergence study on the initial and B-spline surface optimized geometries, with angle of attack adjusted to give $C_L = 0.5$. The difference between the fine and superfine grid levels is about 2 drag counts for the initial geometry, and less than a count for the optimized geometry. Optimization is performed on the coarsest grid level.

Table 10: **Case 4 - Grid parameters for CRM wing grid study**

Grid	Nodes	Off-wall Spacing	y^+
Coarse	925,888	1.5×10^{-6}	0.33
Fine	7,407,104	8.1×10^{-7}	0.17
Superfine	58,456,064	3.9×10^{-7}	0.081

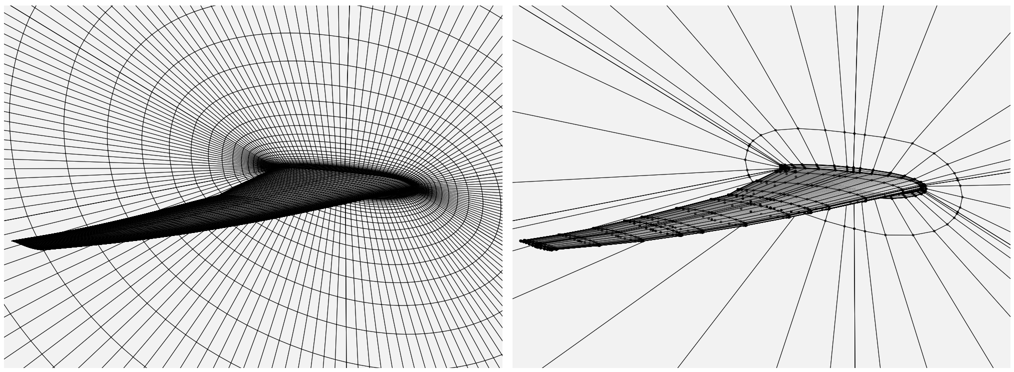

(a) CFD Mesh

(b) B-spline Surface

Figure 24: Case 4 - The computational mesh and B-spline surface for the CRM wing geometry

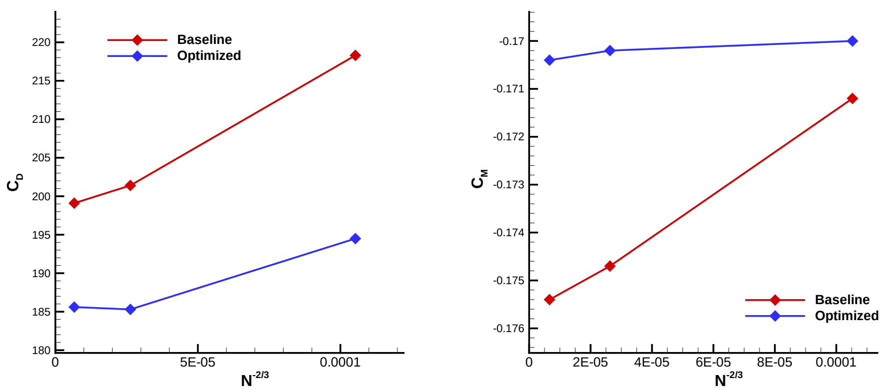

Figure 25: Case 4 - Grid convergence for the CRM wing grid study

Single-Point Optimization Results (Case 4.1)

At each design iteration, the flow solution residual is reduced by 8 orders of magnitude. Figure 26 shows the pressure contours of the baseline and optimized wings using B-splines surface control, and Figure 27 shows the corresponding sectional pressure distributions using both control methodologies, computed on the fine grid level. The wing sections all become thinner except in the root region, which thickens to maintain the initial volume. The leading edge becomes progressively sharper towards the wing tip. The sharp leading edge is likely due to the absence of a low-speed lift constraint for the wing, and could be removed through a geometric constraint if desired. The geometries and pressure distributions obtained using B-spline surface and FFD geometry control are similar, and it is expected that they will become increasingly similar if the optimizations are run longer. The spanwise lift distributions of the initial, FFD, and B-spline optimized geometries evaluated on the fine mesh are compared to the elliptical distribution in Figure 28. While the differences between the geometries optimized using the B-spline and FFD methods are reflected in noticeable differences in drag, as displayed in Table 11, the drag discrepancy gets smaller with refinement. The FFD optimization was also conducted using a “medium” mesh refined by a factor of 1.587 in each direction, giving

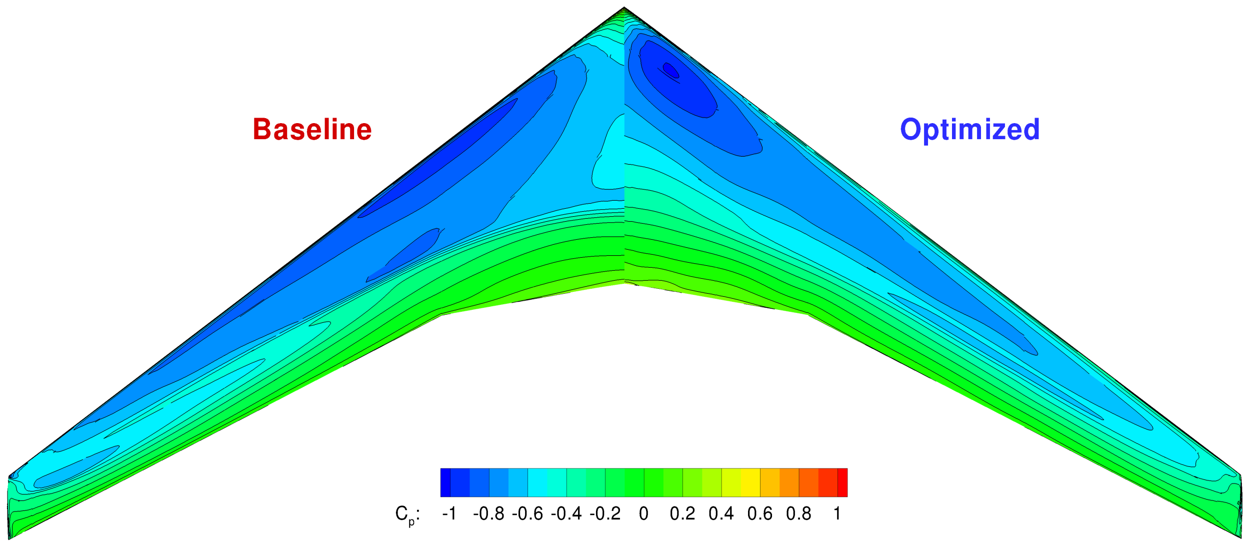

Figure 26: Case 4.1 - Pressure contours for baseline and optimized CRM wing using B-spline surface control

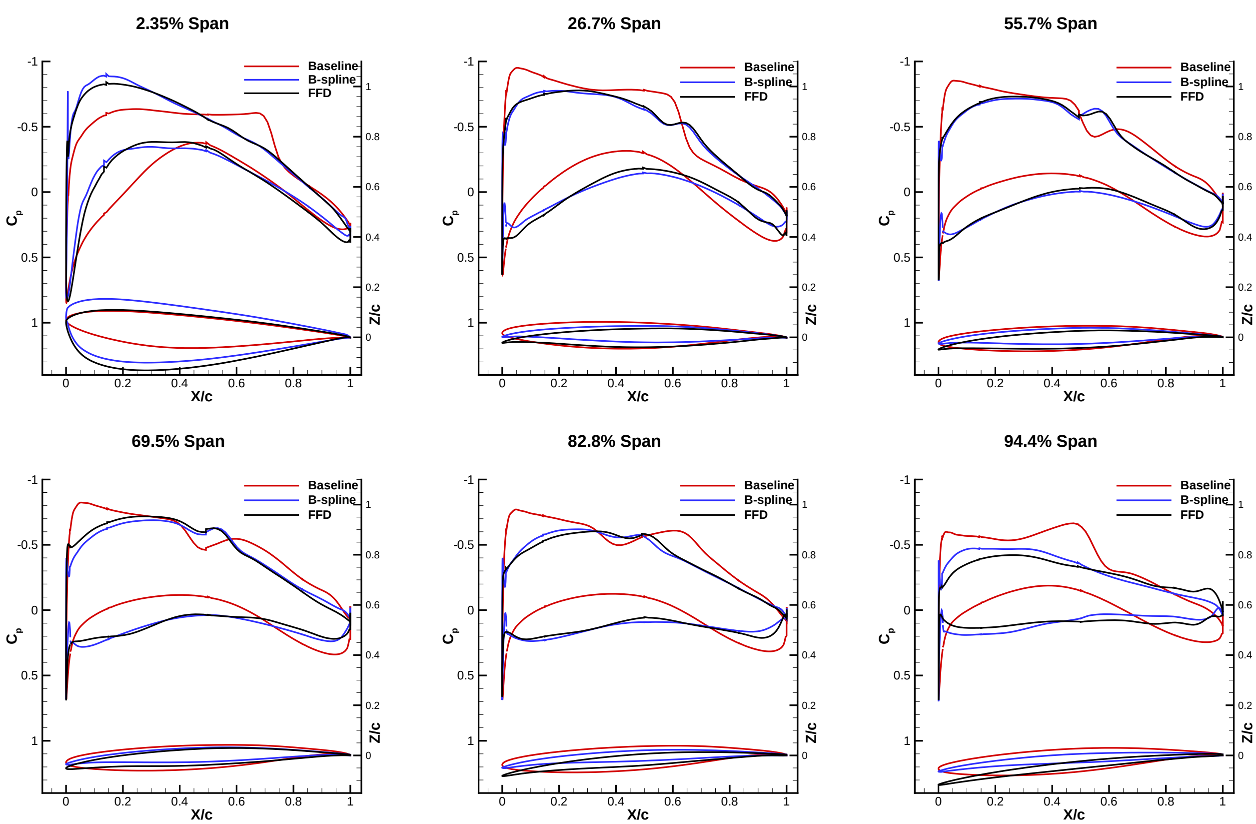

Figure 27: Case 4.1 - Sectional pressure plots and sections for baseline and optimized CRM wings computed on fine mesh

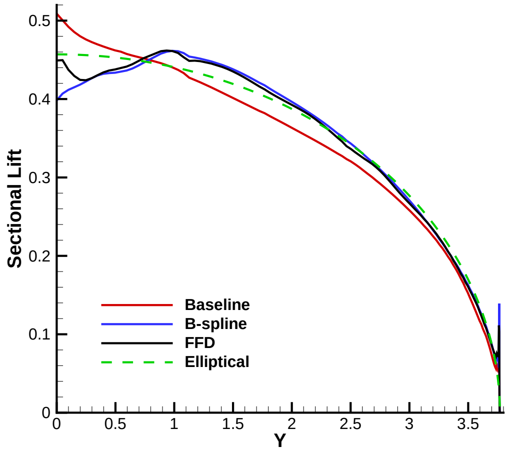

Figure 28: Case 4.1 - Initial, B-spline surface optimized, and elliptical lift distributions

Table 11: Case 4.1 - Results for CRM wing single-point optimization

	Optimization Mesh		Fine Mesh		Superfine Mesh	
	C_D (counts)	C_M	C_D (counts)	C_M	C_D (counts)	C_M
Baseline	218.3	-0.1712	201.5	-0.1747	199.1	-0.1754
B-spline surface control	194.5	-0.1700	185.2	-0.1702	185.6	-0.1704
FFD control	199.0	-0.1700	186.8	-0.1707	185.4	-0.1706
FFD control (medium mesh)	187.1	-0.1700	184.7	-0.1700	183.4	-0.1699

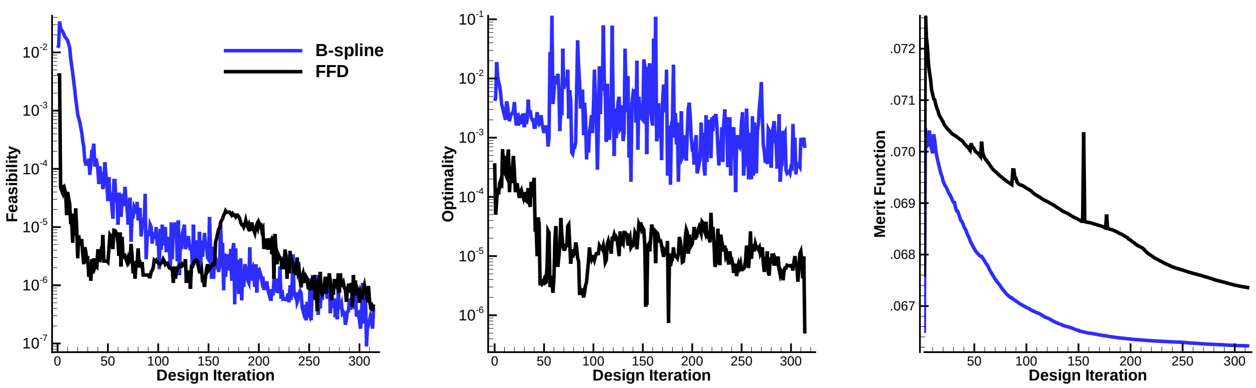

Figure 29: Case 4.1 - Optimization convergence for single-point optimizations of the CRM wing

four times the number of nodes as the coarse mesh, and the results are also included in Table 11. Fine-mesh analysis shows that, as expected, there is some benefit to optimizing on a finer mesh. Due to time limitations, however, the subsequent multi-point optimizations are conducted using the coarse mesh.

Figure 29 shows the SNOPT convergence history for the two optimizations on the coarse mesh. In addition, the convergence of the force and moment coefficients is displayed for the B-spline control case in Figure 30. The feasibility and optimality tolerances are both set at 1×10^{-6}. The optimality measure is reduced by roughly two orders of magnitude relative to its highest value. The difference in convergence rate between B-spline surface and FFD control may be due to design variable scaling. While greater reduction in optimality is possible with further optimization, the merit function plot shows that most of the drag reduction has already been achieved.

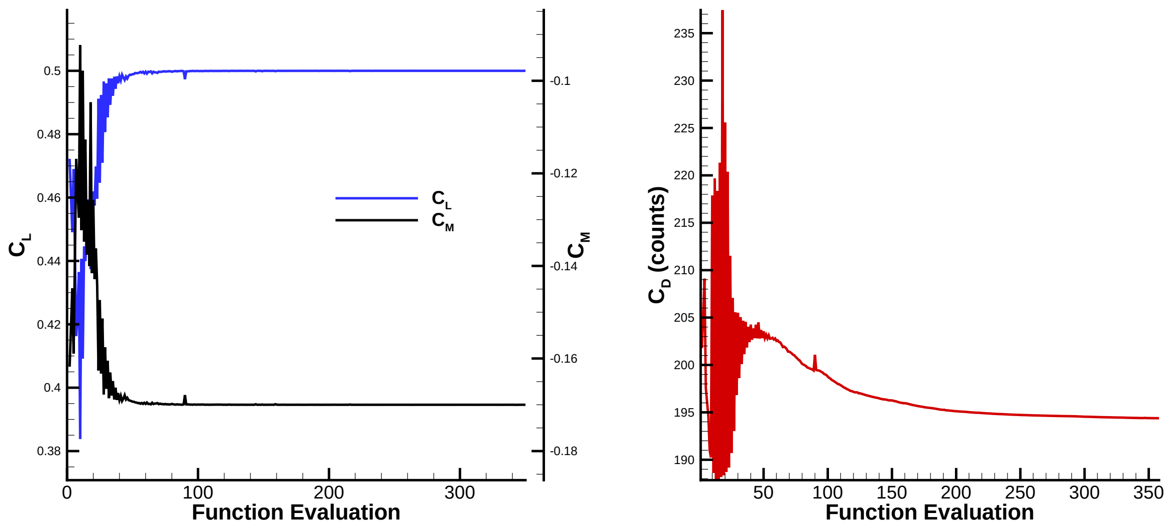

Figure 30: Case 4.1 - Lift, drag, and moment coefficient convergence using B-spline surface control

Multi-point Optimization Results (Cases 4.2 - 4.6)

The degrees of freedom and geometry for the CRM wing multi-point optimization are the same as the single-point problem, with the exception that the angle of attack at each design point is its own design variable. Only B-spline surface control is used. In each design iteration of the multi-point optimization, a flow solution is computed at each of the operating points in parallel. The objective function and gradient are computed as a weighted sum of the results from each of the converged flows. The pitching moment constraint is only satisfied at the nominal design point, which is given the greatest weight $\mathcal{T}_i$. There are four three-point cases: one with variable C_L and constant Mach number, two with variable Mach number and constant C_L, and one with variable Mach number and constant lift. In addition, there is a nine-point case over a range of Mach numbers and lift coefficients. The operating points for each case are summarized in Table 12. The optimization problem is posed as

$$\begin{aligned}
\text{minimize} \; & \sum_{i=1}^{n} \mathcal{T}_i C_{Di} \\
\text{wrt} \; & z, \alpha_i \\
\text{subject to} \; & C_L = C_{Li} \\
& C_M \geq -0.17 \text{ (at nominal design point)} \\
& V \geq V_{\text{baseline}} \\
& z \geq 0.25 \times z_{\text{baseline}}.
\end{aligned}$$

All of the multi-point optimizations are carried out on the coarse mesh, while all the drag polars and analyses are performed with the optimized geometry on the fine level mesh. The optimization histories of cases 4.2 and 4.6 are compared in Figure 31. The three-point optimizations converged more successfully than the nine-point case, as the nine-point case suffered from more flow solver convergence difficulties at the more demanding flow conditions. The drag and moment coefficients of the optimized geometries computed on the fine mesh at the nominal condition are displayed in Table 13. The sections and sectional pressure plots computed on the fine mesh at the nominal condition are displayed in the Appendix, in Figures A.1 and A.2. Figure 32 shows the lift curve and moment curve for the initial, single-point, and multi-point optimized geometries. Figure 33 shows the drag coefficient vs. angle of attack and the drag polar at Mach 0.85. The drag polar shows that all the optimized geometries produce a noticeable improvement in L/D compared

Table 12: Case 4 - Multi-point problem operating points

Case	Point i	Weight $\mathcal{T}_i$	M	C_L	Re
4.2	1	1	0.85	0.450	5.00×10^6
	2	2	0.85	0.500	5.00×10^6
	3	1	0.85	0.550	5.00×10^6
4.3	1	1	0.84	0.500	5.00×10^6
	2	2	0.85	0.500	5.00×10^6
	3	1	0.86	0.500	5.00×10^6
4.4	1	1	0.82	0.500	5.18×10^6
	2	2	0.85	0.500	5.00×10^6
	3	1	0.88	0.500	4.83×10^6
4.5	1	1	0.82	0.537	4.82×10^6
	2	2	0.85	0.500	5.00×10^6
	3	1	0.88	0.466	5.18×10^6
4.6	1	1	0.82	0.483	4.82×10^6
	2	2	0.82	0.537	4.82×10^6
	3	1	0.82	0.591	4.82×10^6
	4	2	0.85	0.450	5.00×10^6
	5	4	0.85	0.550	5.00×10^6
	6	2	0.85	0.550	5.00×10^6
	7	1	0.88	0.442	5.18×10^6
	8	2	0.88	0.466	5.18×10^6
	9	1	0.88	0.513	5.18×10^6

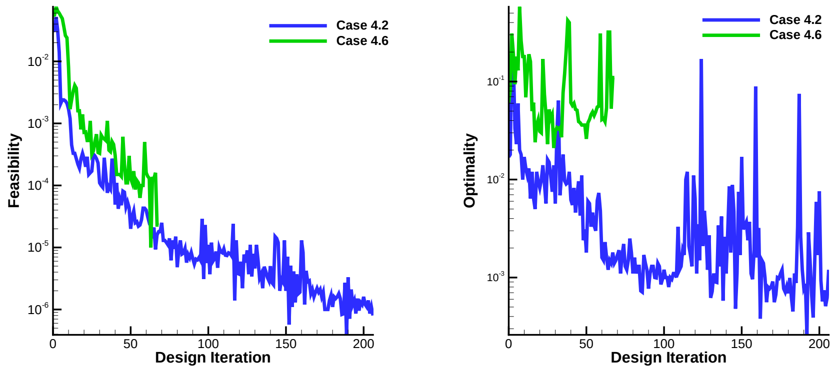

Figure 31: Case 4 - Optimization convergence for multi-point optimizations of the CRM wing

Table 13: Case 4 - Drag counts at nominal operating point $C_L = 0.5$ and Mach 0.85 computed on fine mesh

	C_D (counts)	C_M
Baseline	201.5	-0.1746
Case 4.1	185.2	-0.1702
Case 4.2	185.8	-0.1704
Case 4.3	185.8	-0.1705
Case 4.4	187.8	-0.1711
Case 4.5	187.0	-0.1709
Case 4.6	189.7	-0.1717

the baseline geometry. Compared to the multi-point optimizations, the single-point result shows poorer performance at lower lift coefficients and a slight improvement at the nominal flight condition $C_L = 0.5$. Cases 4.2 and 4.3 perform the best over the range of lift coefficients at $M = 0.85$, which is not surprising as these cases have a narrower range of Mach numbers around the nominal condition. Figure 34a better illustrates the advantage of multi-point optimization over single-point. The single-point geometry shows a higher drag over most Mach numbers, with a slight benefit at the nominal Mach number 0.85. Once again, cases 4.3 and 4.2 perform the best around the operating condition, with the case 4.3 geometry showing slightly better drag. Case 4.4, which was optimized at Mach 0.82 and 0.88 in addition to the nominal Mach number, shows significantly better drag at higher Mach numbers. Figure 34b shows case 4.5 outperforming 4.6, which is expected since case 4.6 was optimized with consideration of additional operating conditions to those of case 4.5.

Figure 35 shows the lift curve and moment curve for the initial and nine-point optimized geometries, at varying Mach numbers. Figure 36 shows the drag coefficient vs. angle of attack and the drag polar at varying Mach numbers. The drag polar shows that for a given Mach number, the drag reduction relative to the baseline curve improves at increased C_L. The drag reduction is marginal for Mach 0.82, but much more significant at Mach 0.85 and 0.88. Figure 37 plots drag coefficient against Mach number for three fixed lifts given by $C_L = 0.45$, 0.50, and 0.55 at Mach 0.85. Again, the drag reduction relative to the baseline geometry increases at higher Mach numbers. The drag and moment coefficients computed on the fine mesh for optimized geometries at their design conditions are summarized in the Appendix in Table A.1.

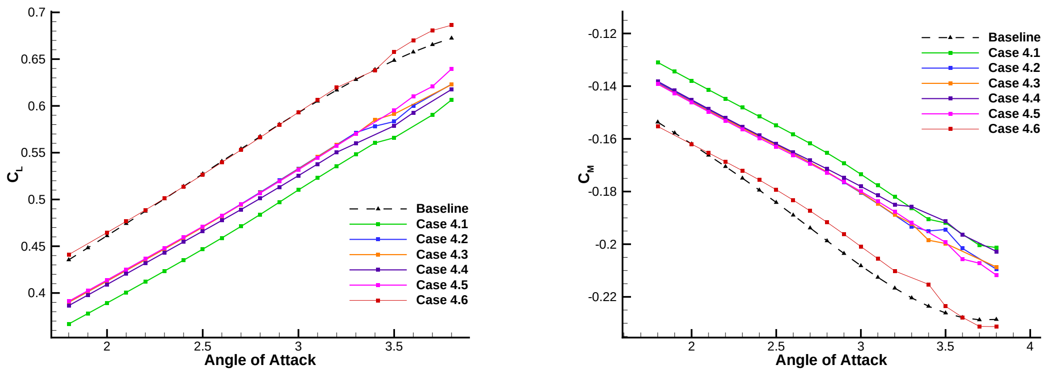

(a) C_L vs. α

(b) C_M vs. α

Figure 32: Case 4 - C_L and C_M vs. α for single-point, three-point, and nine-point optimizations at $M = 0.85$

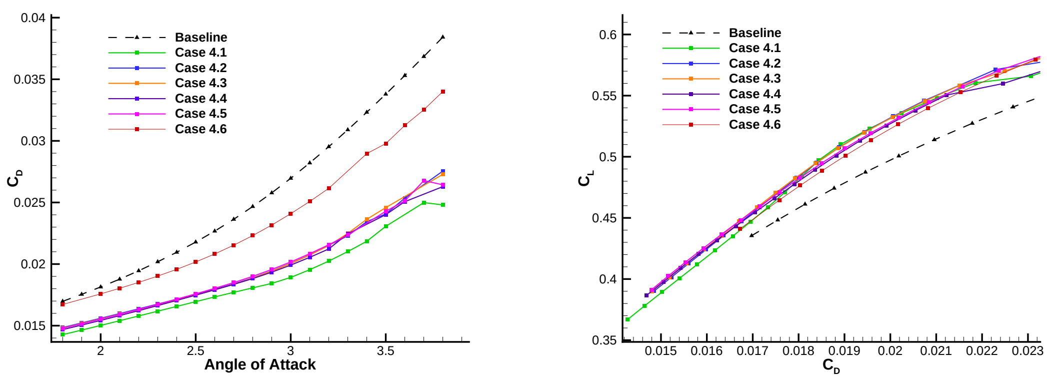

(a) C_D vs. α

(b) Drag polar

Figure 33: Case 4 - C_D vs. alpha and vs. C_L for single-point, three-point, and nine-point optimizations at $M = 0.85$

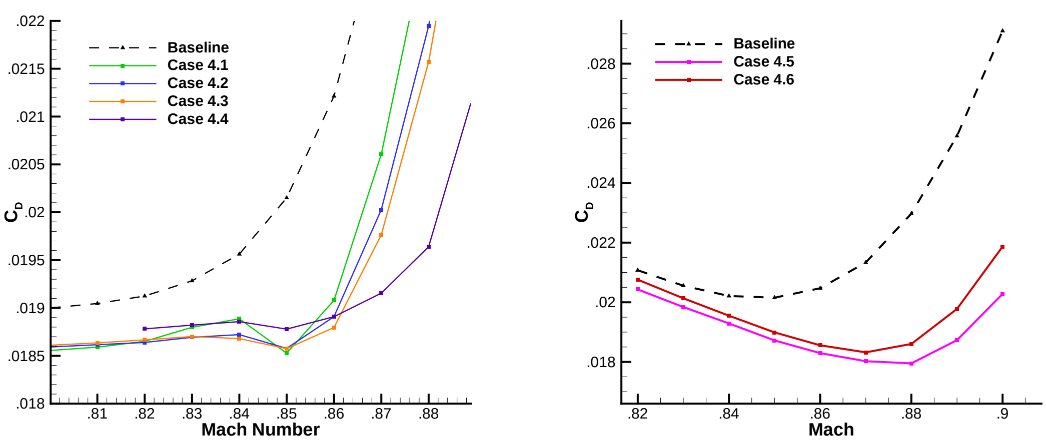

(a) Constant $C_L = 0.5$

(b) Constant lift

Figure 34: Case 4 - C_D vs. Mach number for (a) C_L = 0.5 for Cases 4.1, 4.2, 4.3, 4.4, and (b) constant lift based on C_L = 0.5 and $M = 0.85$ for Cases 4.5 and 4.6

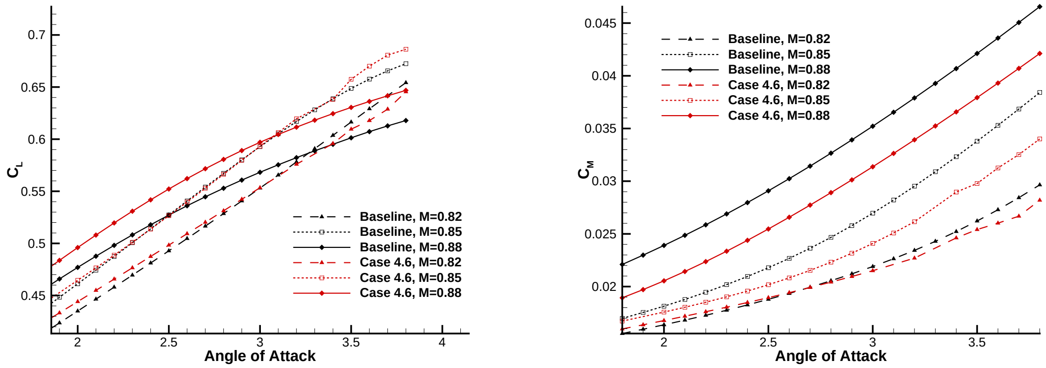

(a) C_L vs. α (b) C_M vs. α

Figure 35: Case 4.6 - C_L and C_M vs. α for nine-point optimization at various Mach numbers

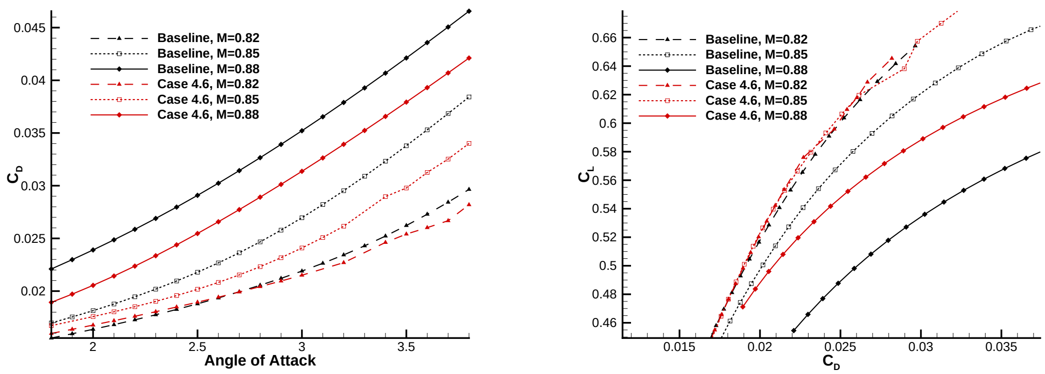

(a) C_D vs. α (b) Drag polar

Figure 36: Case 4.6 - C_D vs. α and vs. C_L for nine-point optimization at various Mach numbers

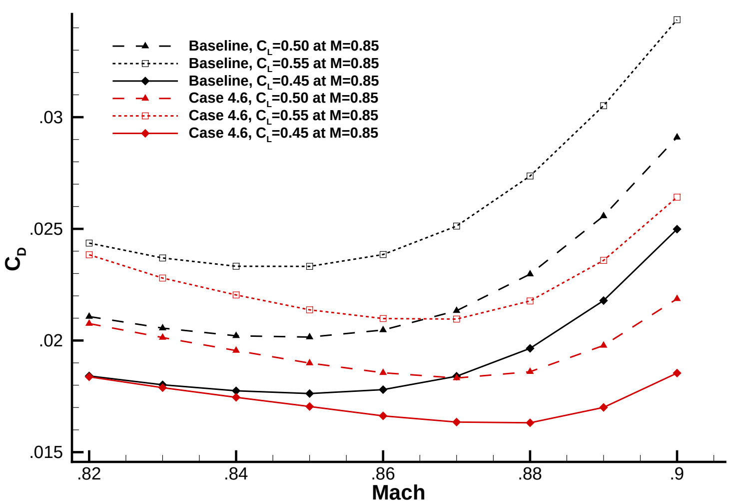

Figure 37: Case 4.6 - C_D vs. Mach number for nine-point optimization at various lifts

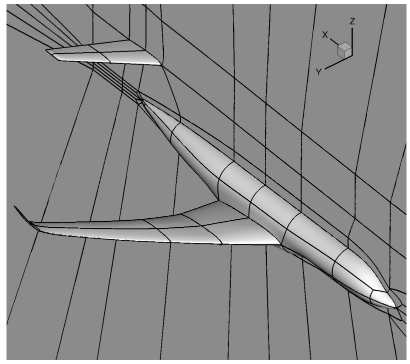

Figure 38: Case A1 - T-tail aircraft geometry and CFD grid blocking topology

IV. Additional Cases

A. Additional Case 1: Wing-Fuselage-Tail Aircraft Optimization

Optimization Problem

The goal of this optimization problem is to determine the optimal trimmed aircraft configuration for minimum drag at a given operating condition where the aircraft configuration includes the wing, fuselage, and a horizontal tail. The operating condition is at Mach number 0.82, $C_L = 0.513$, and Reynolds number 19.1×10^6 based on the wing mean aerodynamic chord. Aerodynamic and geometric constraints imposed on the optimization are as follows. The optimized aircraft geometry must achieve a total lift corresponding to the aircraft weight. In addition, a prescribed spanwise lift distribution is enforced on the wing. The idea behind this approach is that the spanwise lift distribution comes from a medium-fidelity multi-disciplinary optimization, for example based on a panel method for aerodynamic forces and moments. The high-fidelity aerodynamic shape optimization can then be used to eliminate shocks, but should maintain the spanwise lift distribution such that the wing weight from the medium-fidelity optimization remains accurate. A detailed description of the prescribed spanwise lift distribution constraint formulation is given in Osusky et al.[34] To satisfy the trim constraint, pitching moments summed about the aircraft center of gravity, which is at 48.7 feet from the nose of the fuselage, must equal zero. Geometric constraints are imposed on wing volume, sectional areas, and sectional thicknesses to ensure the wing outer mold line is sufficient to accommodate the internal wing structure and fuel volume. Sectional areas and wing volume must be greater than or equal to their initial values. Sectional t/c values must not decrease by more that 25% of their initial values. The optimizer can only vary the wing geometry while keeping the fuselage and tail geometry unchanged. Specifically, the optimizer can only vary wing sectional shapes and twist while maintaining a fixed planform. In addition to the above geometric flexibility, wing and tail angle of incidence with respect to aircraft longitudinal axis, and aircraft angle of attack are allowed to vary. Wing and tail angles of incidence may vary by $\pm 5°$ and $\pm 10°$ respectively. The aircraft angle of attack may vary by $\pm 2°$. The total number of design variables is 508.

Initial Geometry

The initial geometry is a T-tail aircraft configuration shown in Figure 38. RAE 2822 airfoil sections are used for the wing. NASA SC(2)-0012 airfoil sections are used for the horizontal tail, and the vertical fin is not included. The tail geometry is fixed, other than its angle of attack.

Grid

The grid used for the optimization has an average off-wall spacing of 2.5×10^{-6} reference lengths. The reference length is the wing mean aerodynamic chord, which has a value of 12.8 feet. The mesh consists of 7.5×10^6 nodes partitioned over 620 blocks, each of size $23 \times 23 \times 23$ nodes. Figure 38 illustrates the blocking topology.

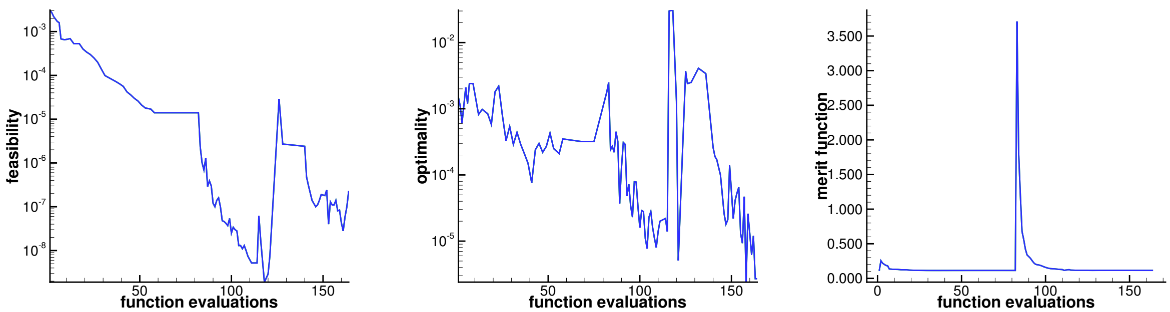

(a) Feasibility (b) Optimality (c) Merit function

Figure 39: Case A1 - Convergence history for the wing-body-tail optimization case

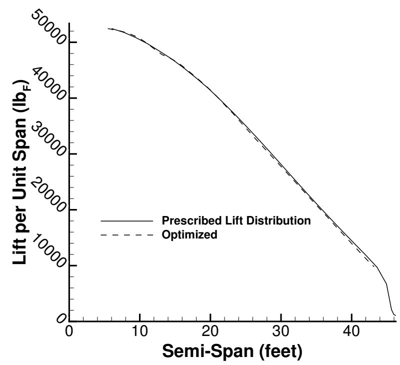

Figure 40: Case A1 - Spanwise lift distribution of optimized wing compared to prescribed lift distribution

Optimization Results

After 164 design iterations, a converged optimization result has been achieved. The optimization convergence history shown in Figure 39 indicates that optimality has been reduced by two orders of magnitude and all aerodynamic and geometric constraints have been satisfied. It can be seen that the optimizer struggles to reduce feasibility until the 83rd design iteration. At this iteration, the constraints on total aircraft lift and wing lift distribution are not satisfied. The aircraft angle of attack has been stuck at its upper bound of 2° since the 4th function evaluation. At the 84th design iteration the optimizer performs an internal reset of its Hessian approximation, which has the desired effect of breaking the feasibility stagnation. The optimizer achieves satisfaction of the total aircraft lift and the wing lift distribution constraints at design iteration 95 by increasing outboard wing twist and reducing aircraft angle of attack to 1.8°. Satisfaction of the spanwise lift distribution constraint is shown in Figure 40, where it can be seen that the optimized lift distribution closely matches the prescribed lift distribution. Performance for the optimized aircraft configuration is summarized in Table 14. The optimized aircraft geometry is evaluated on a 24-million-node fine grid to obtain a more accurate prediction of performance. A comparison of the baseline and optimized aircraft geometries evaluated on the optimization mesh level at the target lift shows an improvement in C_D of 27%. Sections and pressure distributions at several spanwise locations on the initial and optimized wings are shown in Figure 41. The pressure distributions between 42% and 58% span show a shock on the initial wing geometry that has been eliminated on the optimized wing geometry. These results demonstrate the capability of the aerodynamic shape optimization algorithm to design a wing for minimum drag at turbulent, transonic operating conditions for a trimmed T-tail aircraft configuration by eliminating shocks on the wing while maintaining a prescribed spanwise lift distribution. This is one way in which aerodynamic shape optimization can be used to refine a design from a low- or medium-fidelity multidisciplinary optimization.

Table 14: Case A1 - Performance for initial and optimized T-tail aircraft geometries

	Baseline	Optimized	Optimized (Fine Grid Analysis)
C_L	0.513	0.513	0.513
C_D	0.0507	0.0369	0.0309
$\frac{C_L}{C_D}$	10.12	13.89	16.60
C_M	0.000	0.000	0.000
α (°)	3.04	1.86	1.87

(a) 10% Span

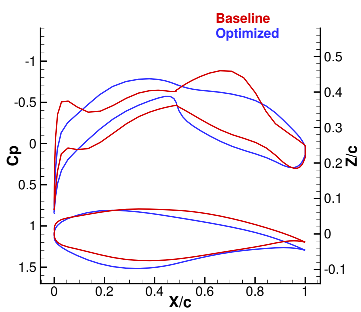

(b) 26% Span

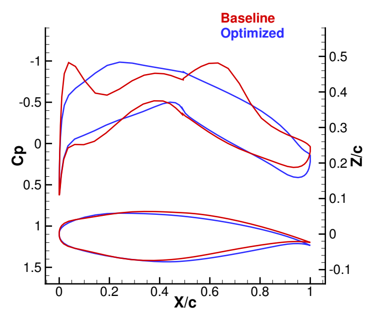

(c) 42% Span

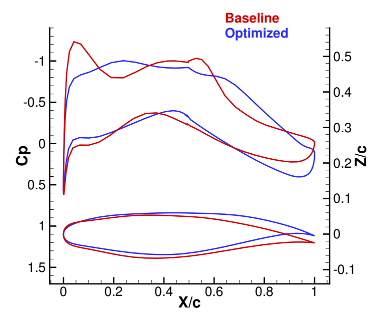

(d) 58% Span

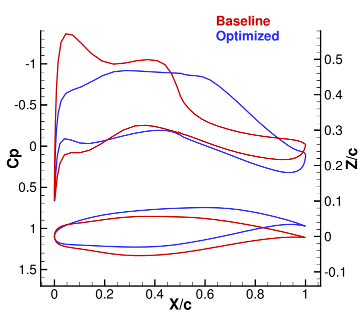

(e) 74% Span

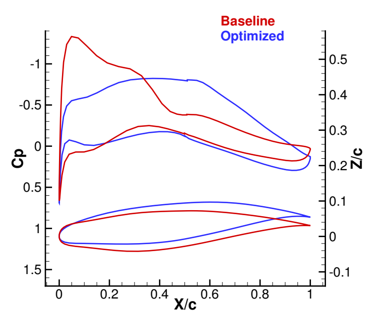

(f) 90% Span

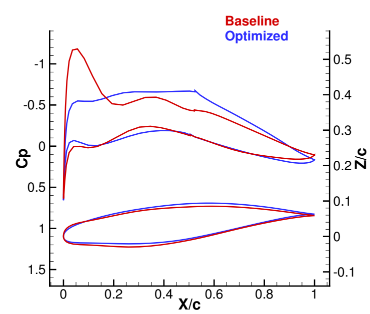

Figure 41: Case A1 - T-tail aircraft wing sections and pressure distributions for initial and optimized geometries

B. Additional Case 2: Box Wing Optimization

Rationale

According to lifting-line theory, the box wing is the lifting system that produces the least amount of induced drag for a given lift, span, and vertical extent.[35] Moreover, the span efficiency of the box wing increases as the height-to-span (h/b) ratio increases; in Prandtl,[35] an equation relating the two quantities is provided:

$$\frac{1}{e} \approx \frac{1 + 0.45(h/b)}{1.04 + 2.81(h/b)}. \tag{9}$$

However, it is unclear how nonlinear aerodynamics impact the span efficiency of the box wing.

The optimal lift distribution of the box wing is generally depicted as the sum of a constant and an elliptical lift distribution that is equally carried by the top and bottom wings, joined at their tips by butterfly-shaped side-force distributions.[36] However, as remarked in Kroo,[37] according to lifting-line theory a vortex loop of constant circulation can be superimposed to such a closed system without changing its total lift and drag. Thus, at least from the perspective of linear theory, for a given h/b ratio there is an infinite number of optimally-loaded box wings.

The intent of this proposed additional case is twofold. First, it is desired to compare the span efficiencies estimated by Eq. (9) with those obtained through numerical optimization when nonlinear aerodynamics are included. The problem is more challenging than in Case 3, since here the flow solver must correctly resolve the flow field induced by the neighboring wings in order for the optimizer to exploit the nonlinearity of the wake. Further, unlike Case 3, the proposed case is not affected by side-edge tip separation, although care must be taken where the wings connect with the tip fins (as discussed below). The second intent of the proposed case is to determine, again through nonlinear aerodynamics, whether the optimal lift distribution of the box wing is unimodal or, as implied in Kroo,[37] multimodal with an infinite number of solutions.

Optimization Problem

Specifically, the objective of this case is to maximize the span efficiency,

$$e = \frac{(L/q_\infty)^2}{\pi b^2 (D/q_\infty)}, \tag{10}$$

of the box wing over a range of h/b ratios by varying twist while constraining lift. The flow is modeled with the Euler equations, and the freestream Mach number is fixed at 0.3. Both the height and span, and thus h/b, are fixed during the course of an optimization. Hence, for a given h/b ratio the optimization problem is

$$\begin{aligned} \text{minimize } & D/q_\infty \\ \text{wrt } & \gamma \\ \text{subject to } & L/q_\infty = (L/q_\infty)_{\text{target}}, \end{aligned}$$

where γ is the twist distribution of the entire system, including the corner transitions as described below. Twist is achieved by true rotation (as opposed to shear) about the leading edge; hence, the planform is not fixed, which is why there is no reference area in the problem formulation. However, the lift target is chosen such that $C_L = 0.5$ at the beginning of the optimization, based on the total planform area. The initial planform area of the top and bottom wing is $S_{1/2} = 12c^2$; thus, $(L/q_\infty)_{\text{target}} = C_L \times S = 0.5 \times 24c^2 = 12c^2$. In practice, only half geometries are considered, so $(L/q_\infty)_{\text{target}} = 6c^2$.

Initial Geometry

A generic box-wing geometry is shown in Figure 42. The positive x, y, and z axes correspond to the chordwise, spanwise, and vertical directions, respectively. The two wings and the tip fin are generated from a sharp NACA 0012 airfoil, which is rotated 90° at the wing extremities to close the system. The purpose of the corner fillets is to reduce compressibility effects. For a given h/b ratio, compressibility effects can be reduced further by increasing the size of h and b relative to c. Here, R (see Figure 42) and b are fixed to $0.15c$ and $12c$, respectively, and h is varied from $1.2c$ to $3.6c$ by increments of $0.6c$. Thus, the five h/b ratios considered in this paper are 0.10, 0.15, 0.20, 0.25, and 0.30. Note that the reference area used in the computation of the initial lift coefficient is $S = 2bc$. As explained above, this leads to $(L/q_\infty)_{\text{target}} = 6c^2$ for all five half geometries. Also note that the arc length of the corner fillets are excluded from the definition of the normalized semi-span (η) and vertical (η_{V}) axes.

Grid

The optimization grids are composed of 42 blocks and 2,569,427 nodes with off-wall spacings of about $10^{-3}c$. All five grids have the same hyperbolic mesh law parameters along the same respective block edges. The angles of attack necessary to achieve the target lift of $6c^2$ at the beginning of the optimizations are 6.1041°, 5.6455°, 5.4058°, 5.2592°, and 5.1586° for $h/b = 0.10$, 0.15, 0.20, 0.25, and 0.30, respectively. As in the previous cases, grid-converged lift and drag values are obtained by performing post-optimization flow solves on superfine grids, here composed of 2154 blocks and 89,560,035 nodes.

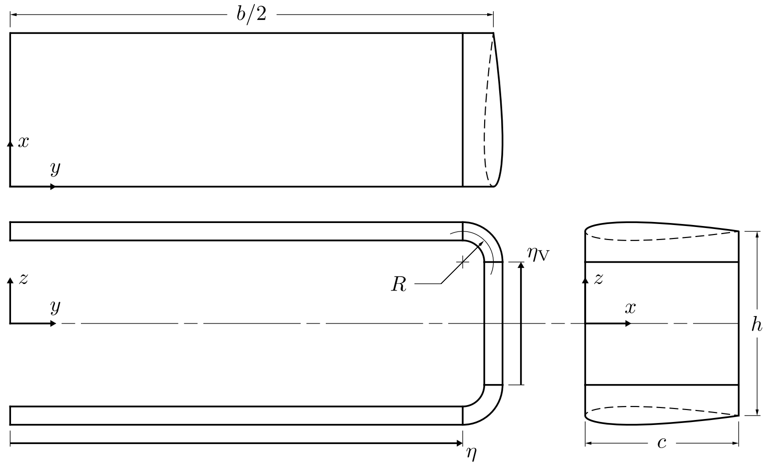

Figure 42: Case A2 - Generic box-wing geometry

Optimization Results

A total of 22 FFD design variables are used to achieve a continuous, piecewise-cubic twist distribution. Twist is applied about the leading edge; thus, the rotation planes are normal to it. For example, the rotation planes for the tip fin are normal to the global xy plane. For the corner fillets, the rotation planes are derived from a linear combination of xz and xy planes.

The top and bottom wings are each assigned 7 design variables that are evenly distributed along $\eta \in [0, 1]$; similarly, the tip fin is assigned 4 design variables that are evenly distributed along $\eta_V \in [0, 1]$. Each corner fillet has 2 additional design variables that are evenly spaced between its 2 tip design variables. However, in order to prevent the development of overly wavy surfaces at the corner fillets, these additional design variables are constrained to linearly interpolate the tip design variables; hence, the optimization problems have effectively 18 design variables each.

All five optimizations converged to optimality and feasibility tolerances of 1×10^{-6} and 1×10^{-7}, respectively, in 28 major SNOPT iterations or less. The convergence histories are very similar to those shown in Figure 21.

As seen from Figure 43, the (inverse of the) span efficiencies obtained through the Euler-based optimizations are in good agreement with those estimated by linear theory, i.e. Eq. (9). The discrepancies can be attributed to the definition of h and b in Figure 42. If the true bounding boxes of the overall systems are used instead, i.e. when airfoil thickness is accounted for, the h/b ratios are in fact slightly larger. Consequently, with the corrected h/b ratios the red curve shifts to the right and falls almost exactly on the black curve.

It is worth pointing out that airfoils designed for biplanes should have substantially different camber than those designed for monoplanes.[38] Thus, in the future it would be interesting to include sectional design variables in the optimization problems. Preliminary results suggest that, at least for the $h/b = 0.20$ case, a 5 to 10% higher span efficiency than predicted by Eq. (9) is possible,[39] even with the more conservative bounding box definition of h/b.

The sectional force coefficients along the spanwise (vertical) axis of the wings (tip fin) are plotted in Figure 44. Note that the force vectors used to compute these coefficients are not oriented according to the angle of attack, but rather according to the Cartesian axis normal to each surface, i.e. the global z axis for the wings and the global y axis for the tip fin. Unlike the span efficiencies, the vertical force distributions differ significantly compared to those typically depicted.[36] Here the bottom wing carries significantly more lift than the top wing, although as the h/b ratio increases, the load is progressively shifted to the top wing. The side-force distribution adapts to this shift while remaining relatively similar in shape. Finally, the force distribution of the top wing appears to be more elliptical than that of the bottom wing; in general, it is also smoother.

As a check for multimodality, the $h/b = 0.20$ case was repeated five times, each time with a different starting twist distribution. Specifically, each case was started from a separate set of randomly-generated design variables ranging from -10° to 10°. All cases converged to the same solution (plus or minus numerical tolerances), suggesting that for a given h/b ratio the optimal force distribution of the box wing, based on the Euler equations, is unimodal.

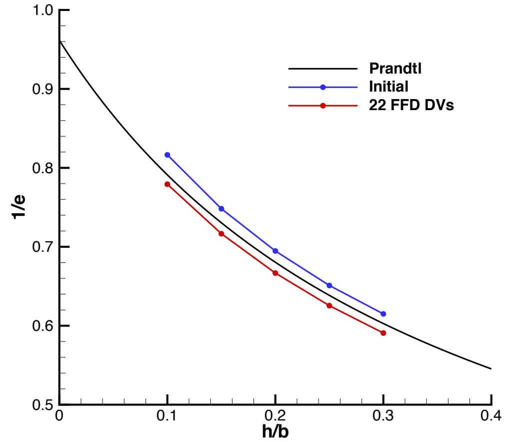

Figure 43: Case A2 - Inverse of span efficiency versus height-to-span ratio

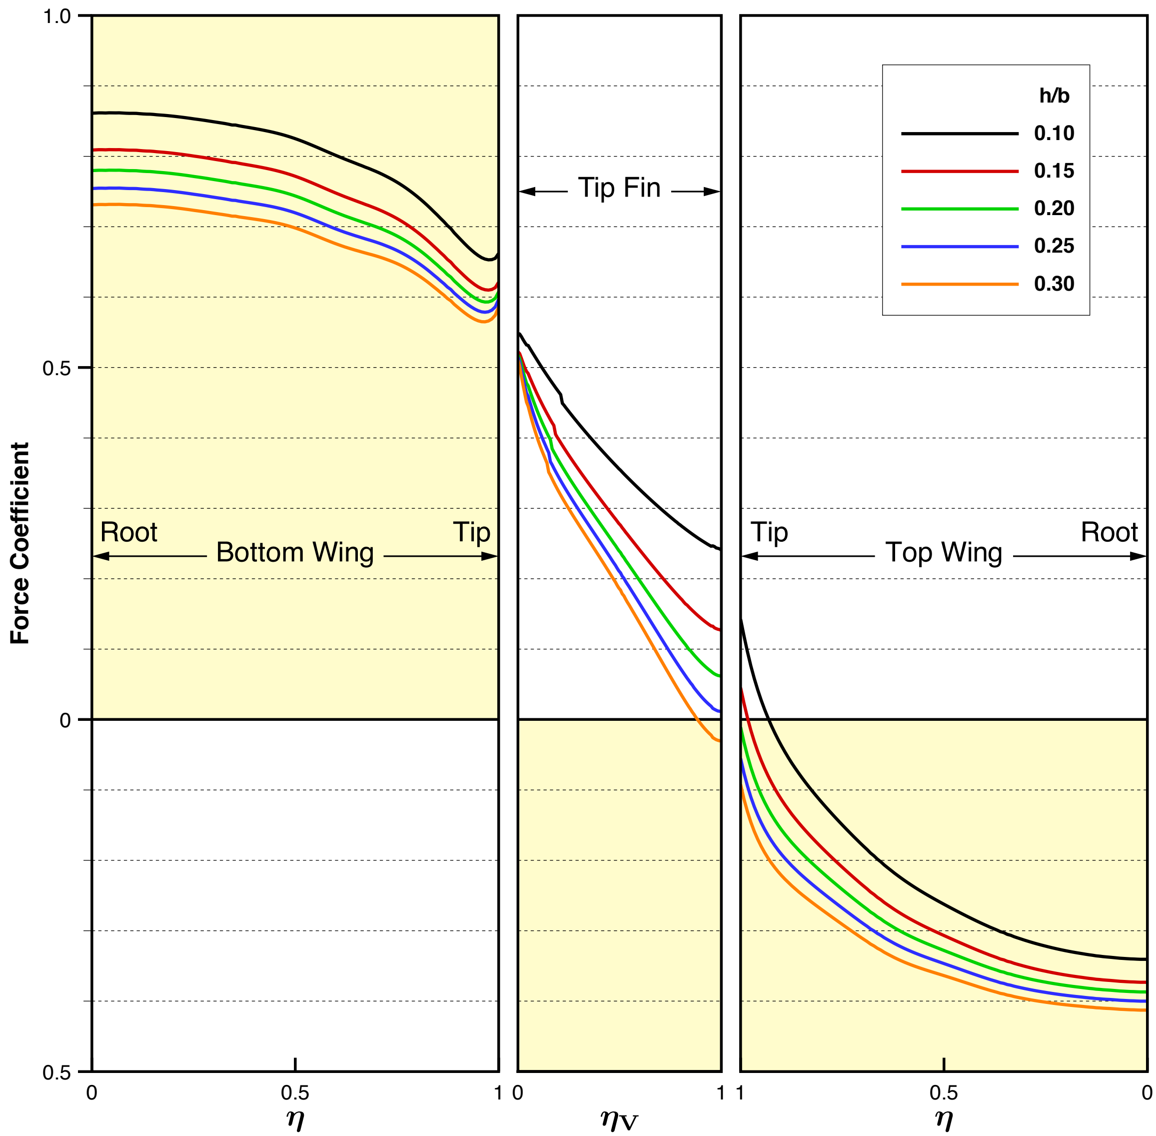

Figure 44: Case A2 - Optimal force coefficient distributions for five different height-to-span ratios. The regions with a yellow background correspond to positive force coefficients (refer to Figure 42).

V. Conclusions

The Jetstream aerodynamic shape optimization algorithm is applied to four benchmark optimization problems as part of the Aerodynamic Design Optimization Discussion Group. For the NACA 0012 optimization, the shock is weakened and pushed downstream by thickening the airfoil at the leading and trailing edges, reducing C_d to 42.2 counts in the best case. For the RAE 2822 case, successful optimizations eliminate the shock, reducing C_d to 119.2 counts in the best case. Both two-dimensional cases yield non-unique solutions that hinder optimization convergence. The twist optimizations converged successfully, giving nearly elliptical lift distributions and span efficiency factors very close to unity. The single-point CRM wing optimization on the medium mesh with FFD control reduces C_D to 183.4 counts. The multi-point optimizations give higher drag at the nominal condition, but lower drag at other Mach number and lift conditions, especially at higher Mach numbers. In all of the CRM wing cases, significant shape changes and performance improvements are achieved.

In addition, two additional cases are proposed as candidates to be added to the benchmark problem suite. The first is a lift-constrained drag minimization with a prescribed spanwise lift distribution for a wing-fuselage-tail geometry. The second is a box-wing, which allows for the optimization of an unconventional shape, while still providing a basis for comparison to lifting-line theory.

Appendix: Additional Figures and Table for CRM Wing

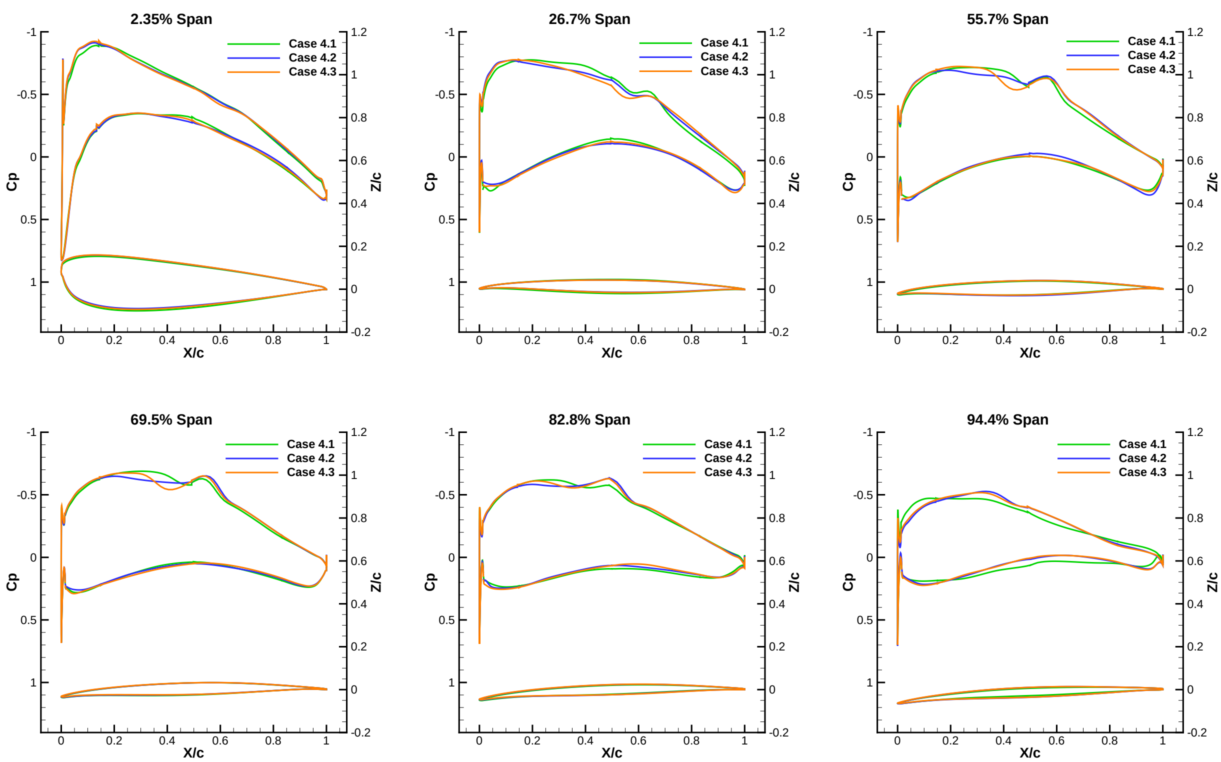

Figure A.1: Case 4 - Sectional pressure plots and sections for optimized CRM wings, computed on the fine mesh at the nominal condition

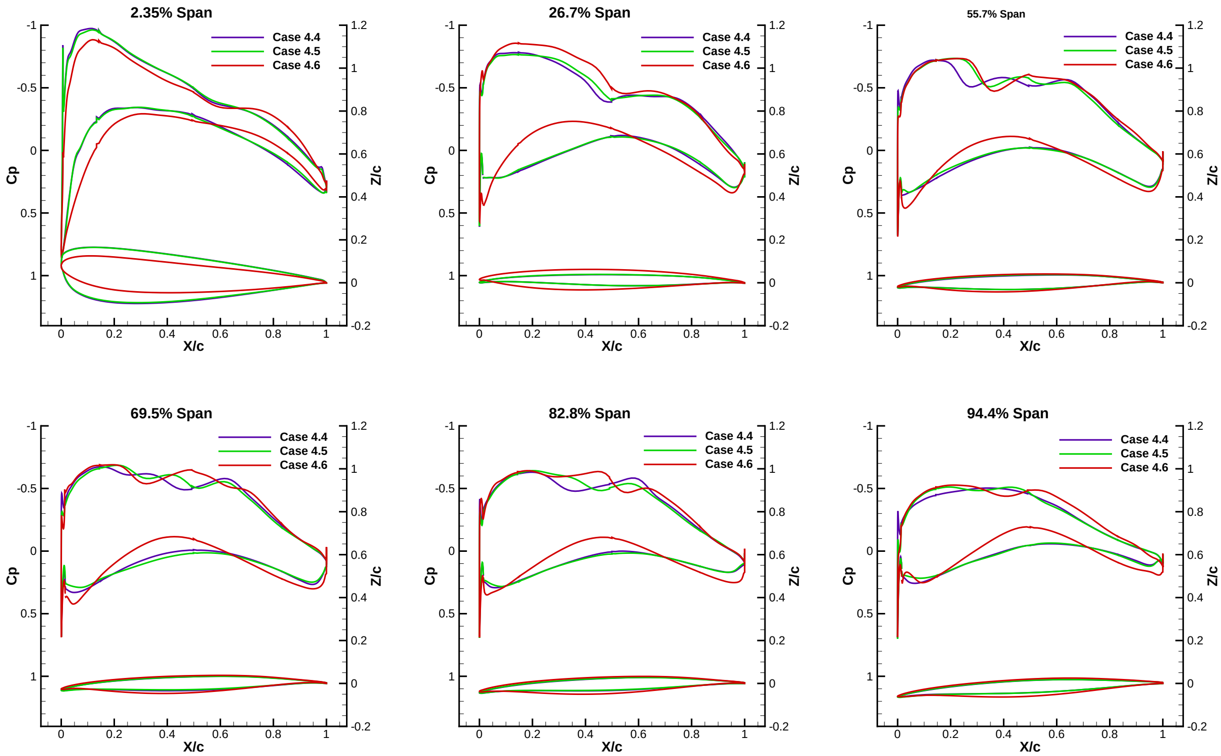

Figure A.2: Case 4 - Sectional pressure plots and sections for optimized CRM wings, computed on the fine mesh at the nominal condition

Acknowledgments

The authors gratefully acknowledge the financial assistance from the National Sciences and Engineering Research Council, the Ontario Graduate Scholarship program, Bombardier, the Canada Research Chairs program, and the University of Toronto. Computations were performed on the GPC supercomputer at the SciNet HPC Consortium, part of Compute Canada. 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.

Table A.1: Case 4 - Summary of multi-point force coefficients for baseline and optimized geometries computed on the fine mesh

				Baseline		Optimized	
Case	Point	M	C_L	C_D	C_M	C_D	C_M
4.2	1	0.85	0.450	176.2	-0.1582	168.5	-0.1564
	2	0.85	0.500	201.5	-0.1747	185.8	-0.1704
	3	0.85	0.550	233.4	-0.1923	209.7	-0.1861
4.3	1	0.84	0.500	195.6	-0.1713	186.8	-0.1697
	2	0.85	0.500	201.5	-0.1747	185.7	-0.1705
	3	0.86	0.500	212.1	-0.1801	188.0	-0.1737
4.4	1	0.82	0.500	191.2	-0.1680	187.8	-0.1681
	2	0.85	0.500	201.5	-0.1747	187.8	-0.1711
	3	0.88	0.500	260.3	-0.1902	196.4	-0.1831
4.5	1	0.82	0.537	210.7	-0.1777	204.3	-0.1779
	2	0.85	0.500	201.5	-0.1747	187.0	-0.1708
	3	0.88	0.466	229.7	-0.1794	179.5	-0.1709
4.6	1	0.82	0.483	184.1	-0.1632	183.7	-0.1637
	2	0.82	0.537	210.7	-0.1777	207.7	-0.1783
	3	0.82	0.591	243.6	-0.1925	239.4	-0.1917
	4	0.85	0.450	176.3	-0.1584	170.4	-0.1579
	5	0.85	0.500	201.5	-0.1747	189.8	-0.1718
	6	0.85	0.550	233.2	-0.1923	213.7	-0.1865
	7	0.88	0.442	196.5	-0.1623	163.2	-0.1564
	8	0.88	0.466	229.7	-0.1794	186.0	-0.1742
	9	0.88	0.513	273.6	-0.1947	217.7	-0.1924

References

[1]IATA, "IATA Annual Review 2013," 2013.

[2]IATA, "Fact Sheet: Industry Statistics," 2012.

[3]IATA, "IATA Annual Review 2014," 2014.

[4]Levy, D., Laflin, K., Tinoco, E., Vassberg, J., Mani, M., Rider, B., Rumsey, C., Wahls, R., Morrison, J., Brodersen, O., Crippa, S., Mavriplis, D., and Murayama, M., "Summary of Data from the Fifth AIAA CFD Drag Prediction Workshop," *51st AIAA Aerospace Sciences Meeting*, Dallas, Texas, January 2013.

[5]Telidetzki, K., Osusky, L., and Zingg, D. W., "Application of Jetstream to a Suite of Aerodynamic Shape Optimization Problems," *52nd Aerospace Sciences Meeting*, No. AIAA-2014-0571, National Harbor, Maryland, 2014.

[6]Hicken, J. E. and Zingg, D. W., "Aerodynamic optimization algorithm with integrated geometry parameterization and mesh movement," *AIAA Journal*, Vol. 48, No. 2, 2010, pp. 400–413.

[7]Hicken, J. E., *Efficient Algorithms for Future Aircraft Design: Contributions to Aerodynamic Shape Optimization*, Ph.D. thesis, University of Toronto, 2009.

[8]Gagnon, H. and Zingg, D. W., "Two-Level Free-Form and Axial Deformation for Exploratory Aerodynamic Shape Optimization," *AIAA Journal,* Accepted, November 2014.

[9]Gagnon, H. and Zingg, D. W., "Geometry Generation of Complex Unconventional Aircraft with Application to High-Fidelity Aerodynamic Shape Optimization," *21st AIAA Computational Fluid Dynamics Conference*, No. AIAA-2013-2850, San Diego, California, 2013.

[10]Truong, A. H., Oldfield, C. A., and Zingg, D. W., "Mesh movement for a Discrete-Adjoint Newton-Krylov Algorithm for Aerodynamic Optimization," *AIAA Journal*, Vol. 46, No. 7, 2008, pp. 1695–1704.

[11]Hicken, J. E. and Zingg, D. W., "Parallel Newton-Krylov Solver for the Euler Equations Discretized Using Simultaneous-Approximation Terms," *AIAA Journal*, Vol. 46, No. 11, 2008, pp. 2773–2786.

[12]Osusky, M. and Zingg, D. W., "A parallel Newton-Krylov-Schur flow solver for the Navier-Stokes equations discretized using summation-by-parts operators," *AIAA Journal*, Vol. 51, No. 12, 2013, pp. 2833–2851.

[13]Jameson, A., Schmidt, W., and Turkel, E., "Numerical solutions of the Euler equations by finite volume methods using Runge-Kutta time-stepping schemes," *14th Fluid and Plasma Dynamics Conference*, No. AIAA-81-1259, Palo Alto, California, 1981.

[14]Pulliam, T. H., "Lecture Notes for the von Kármán Inst. for Fluid Dynamics Lecture Series: Numerical Techniques for Viscous Flow Computation in Turbomachinery Bladings," Tech. rep., , Rhode-Saint-Genèse, Belgium, 1986.

[15]Swanson, R. C. and Turkel, E., "On central-difference and upwind schemes," *Journal of Computational Physics*, Vol. 101, No. 2, 1992, pp. 292–306.

[16]Chernukhin, O. and Zingg, D. W., "Multimodality and Global Optimization in Aerodynamic Design," *AIAA Journal*, Vol. 51, No. 6, 2013, pp. 1342–1354.

[17]Nocedal, J. and Wright, S., *Numerical Optimization*, Springer series in operations research and financial engineering, Springer, 1999.

[18]De Sturler, E., "Nested Krylov methods based on GCR," *Journal of Computational and Applied Mathematics*, Vol. 67, No. 1, 1996, pp. 15–41.

[19]De Sturler, E., "Truncation strategies for optimal Krylov subspace methods," *SIAM Journal of Numerical Analysis*, Vol. 36, No. 3, 1999, pp. 864–889.

[20]Hicken, J. E. and Zingg, D. W., "A simplified and flexible variant of GCROT for solving nonsymmetric linear systems," *SIAM Journal on Scientific Computing*, Vol. 32, No. 3, 2010, pp. 1672–1694.

[21]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.

[22]Vassberg, J. C., Harrison, N. A., and Roman, D. L., "A systematic study on the impact of dimensionality for a two-dimensional aerodynamic optimization model problem," *29th AIAA Applied Aerodynamics Conference*, No. AIAA-2011-3176, Honolulu, Hawaii, 2011.

[23]Bisson, F., Nadarajah, S., and Shi-Dong, D., "Adjoint-Based Aerodynamic Optimization of Benchmark Problems," *52nd Aerospace Sciences Meeting*, No. AIAA-2014-0412, National Harbor, Maryland, 2014.

[24]Ou, K., Jameson, A., and Vassberg, J. C., "Airfoils supporting non-unique transonic solutions for unsteady viscous flows," *7th AIAA Theoretical Fluid Mechanics Conference*, No. 2014-2927, Atlanta, GA, 2014.

[25]Jameson, A., Vassberg, J. C., and Ou, K., "Further studies of airfoils supporting non-unique solutions in transonic flow," *AIAA Journal*, Vol. 50, No. 12, 2012, pp. 2865–2881.

[26]Cook, P. H., McDonald, M. A., and Firmin, M. C. P., "Aerofoil RAE 2822 - Pressure Distributions, and Boundary Layer and Wake Measurements," *Experimental Data Base for Computer Program Assessment*, No. 138, AGARD Advisory Report, 1979, pp. A6–1–A6–77.

[27]LeDoux, S. T., Vassberg, J. C., Young, D. P., Fugal, S., Kamenetskiy, D., Huffman, W. P., Melvin, R. G., and Smith, M. F., "A Study Based on the AIAA Aerodynamic Design Optimization Discussion Group Test Cases," *AIAA Journal*, Accepted, 2014.

[28]Nemec, M. and Zingg, D. W., "Newton-Krylov algorithm for aerodynamic design using the Navier-Stokes equations," *AIAA Journal*, Vol. 40, No. 6, 2002, pp. 1146–1154.

[29]Buckley, H. P., Zhou, B. Y., and Zingg, D. W., "Airfoil Optimization Using Practical Aerodynamic Design Requirements," *Journal of Aircraft*, Vol. 47, No. 5, 2010, pp. 1707–1719.

[30]Buckley, H. P. and Zingg, D. W., "Approach to Aerodynamic Design through Numerical Optimization," *AIAA Journal*, Vol. 51, No. 8, 2013, pp. 1972–1981.

[31]Osusky, L. and Zingg, D. W., "Application of an efficient Newton-Krylov algorithm for aerodynamic shape optimization based on the Reynolds-Averaged Navier-Stokes equations," *21st AIAA Computational Fluid Dynamics Conference*, No. AIAA-2013-2584, San Diego, California, 2013.

[32]Reist, T. A. and Zingg, D. W., "Aerodynamic Shape Optimization of a Blended-Wing-Body Regional Transport for a Short Range Mission," *31st AIAA Applied Aerodynamics Conference*, No. AIAA-2013-2414, San Diego, California, 2013.

[33]Gagnon, H. and Zingg, D. W., "High-fidelity Aerodynamic Shape Optimization of Unconventional Aircraft through Axial Deformation," *52nd Aerospace Sciences Meeting*, No. AIAA-2014-0908, National Harbor, Maryland, 2014.

[34]Osusky, L., Buckley, H., and Zingg, D. W., "Drag Minimization Based on the Reynolds-Averaged Navier-Stokes Equations Using a Parallel Newton-Krylov Approach," *Submitted to AIAA Journal*, March 2014.

[35]Prandtl, L., "Induced Drag of Multiplanes," Tech. rep., NACA TN-182, March 1924.

[36]Lange, R. H., Cahill, J. F., Bradley, E. S., Eudaily, R. R., Jenness, C. M., and MacWilkinson, D. G., "Feasibility study of the transonic biplane concept for transport aircraft application," Tech. rep., NASA CR-132462, Lockheed-Georgia Company, June 1974.

[37]Kroo, I., "Drag Due to Lift: Concepts for Prediction and Reduction," *Annual Review of Fluid Mechanics*, Vol. 33, January 2001, pp. 587–617.

[38]Addoms, R. B. and Spaid, F. W., "Aerodynamic Design of High-Performance Biplane Wings," *Journal of Aircraft*, Vol. 12, No. 8, August 1975, pp. 629–630.

[39]Gagnon, H. and Zingg, D. W., "Aerodynamic Optimization Trade Study of a Box-Wing Aircraft Configuration," *56th AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference*, Kissimmee, Florida, January 2015.