RANS-based Aerodynamic Shape Optimization Investigations of the Common Research Model Wing

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

The aerodynamic shape optimization of transonic wings requires Reynolds-averaged Navier–Stokes (RANS) modeling due to the strong nonlinear coupling between airfoil shape, wave drag, and viscous effects. While there has been some research dedicated to RANS-based aerodynamic shape optimization, there has not been an benchmark case for researchers to compare their results. In this investigations, a series of aerodynamic shape optimizations of the Common Research Model wing defined for the Aerodynamic Design Optimization Workshop are presented. The computational fluid dynamics solves Reynolds-averaged Navier–Stokes equations with a Spalart–Allmaras turbulence model. A gradient-based optimization algorithm is used in conjunction with a discrete adjoint method that computes the derivatives of the aerodynamic forces. The drag coefficient at the nominal flight condition is minimized subject to lift, pitching moment and geometric constraints. A multilevel acceleration technique is used to reduce the computational cost. A total of 768 shape design variables are considered, together with a grid with 28.8 million cells. The drag coefficient of the optimized wing is reduced by 8.5% relative to the baseline. The single-point design has a sharp leading edge that is prone to flow separation at off-design conditions. A more robust design is achieved through a multipoint optimization, which achieves more reliable performance when lift coefficient and Mach number are varied about the nominal flight condition. To test the design space for local minima, randomly generated initial geometries are optimized, and a flat design space with multiple local minima was observed.

I.   Introduction

Recent advances in high performance computing have enabled the deployment of full-scale physics-based numerical simulations and optimization in academia and industry. Computational fluid dynamics (CFD) tools and numerical optimization techniques have been widely adopted to shorten design cycle times and to explore design spaces more effectively. High-fidelity methods enable engineers to perform detailed designs earlier in the design process, allowing them to better understand the design trade-offs and to make more informed design decisions. In addition, advances in sensitivity analysis via the adjoint method [1] have dramatically reduced the computational effort required for aerodynamic shape optimization. However, the optimal strategies for solving aerodynamic shape optimization problems are still not obvious. Performing aerodynamic shape optimization on a large grid size remains a challenging task. The researchers from the aerodynamic shape optimization community created four benchmark problems to test aerodynamic optimization methods in a constrained design space, and organized the Aerodynamic Design Optimization Workshop. The benchmark problems range from the optimization of a two-dimensional airfoil using the Euler equations, to three-dimensional shape optimization using the Navier–Stokes. In this paper, we present the results of the most complex benchmark problem among the four test cases: the lift-constrained drag minimization of the Common Research Model (CRM) wing with flow governed by the Reynolds-Averaged Navier–Stokes equations (RANS) [2].

In addition to the results of this optimization study, we developed a multilevel optimization acceleration technique to increase the performance of aerodynamic shape optimization. This method is analogous to the full-multigrid strategy often used in CFD. We perform the consecutive optimizations with coarsened grids to reduce the computational time needed on the finest grid. Using this method, aerodynamic shape optimization with a large grid size requires significantly less computational time.

---

[*]PhD Candidate, Department of Aerospace Engineering, University of Michigan, AIAA Student Member
[†]Postdoc Research Fellow, Department of Aerospace Engineering, University of Michigan, AIAA Student Member
[‡]Associate Professor, Department of Aerospace Engineering, University of Michigan, AIAA Associate Fellow

The majority of the aerodynamic shape optimization problems in the literature are solved with gradient-based optimizers [3, 4, 5, 6, 7]. High-fidelity aerodynamic shape optimization with a large number of design variables has the potential to have multiple local minima. The problem is that due to the high number of dimensions, the design space is difficult to visualize and it is challenging to identify local minima. Several authors addressed the local minima issue with non-gradient-based optimizations [8, 9]. However, in these cases, a smaller number of design variables had to be used due to the high number of function calls required by non-gradient-based methods. Chernukin and Zingg presented a multi-start algorithm and a hybrid optimizer to address the multi-modality issue in aerodynamic optimization [10]. In this paper, we further explore the multi-modality of aerodynamic shape optimization problems by performing the RANS-based shape optimizations starting from randomly generated geometries. The design space is then visualized by taking slices in this multi-dimensional space between different local minima. The results reveal a flat design space near the optimum, where the local minima exhibit nearly, and the corresponding airfoil are similar, but visibly different.

The paper is organized as follows. The numerical tools used in this work are described in Section II. The problem formulation, the mesh, and the baseline geometry are described in Section III. The aerodynamic shape optimization of the CRM wing is presented in Section IV. The multilevel optimization acceleration technique is discussed in Section V, followed by the multi-point optimization in Section VII. The thickness constraint and the number design variables are examined in more detail in Sections VI and VIII. Finally, the multi-modality of the aerodynamic design optimization cases is studied in Section IX, followed by the conclusions.

II.   Methodology

This section describes the numerical tools and methods we used for the optimization studies. These tools are components of the framework for multidisciplinary design optimization (MDO) of Aircraft Configurations with High fidelity (MACH) [11, 12]. MACH can perform the simultaneous optimization of aerodynamic shape and structural sizing variables considering aeroelastic deflections. However, in this paper we focus solely on the aerodynamic shape optimization.

A.   Geometric Parametrization

We use an free-form deformation (FFD) volume approach to parametrize the wing geometry [13]. The FFD volume parametrizes the geometry changes rather than the geometry itself, resulting in a more efficient and compact set of geometry design variables, thus making it easier to manipulate complex geometries. Any geometry may be embedded inside the volume by performing a Newton search to map the parameter space to physical space. All the geometric changes are performed on the outer boundary of the FFD volume. Any modification of this outer boundary can be used to indirectly modify the embedded objects. The key assumption of the FFD approach is that the geometry has constant topology throughout the optimization process, which is usually the case for wing design. In addition, since FFD volumes are tri-variate B-spline volumes, the sensitivities of any point inside the volume can be easily computed. Figure 1 shows the FFD volume and the geometric control points for the aerodynamic shape optimization.


Figure 1. The shape design variables are the z-coordinates of 768 FFD control points.

B.   Mesh Perturbation

Since FFD volumes modify the geometry during the optimization, we must perturb the mesh for the CFD analysis to solve for the modified geometry. The mesh perturbation scheme used in this work is a hybridization of algebraic

and linear elasticity methods [13]. The idea behind the hybrid warping scheme is to apply a linear-elasticity-based warping scheme to a coarse approximation of the mesh to account for large, low-frequency perturbations, and to use the algebraic warping approach to attenuate small, high-frequency perturbations. For the results in this paper, the additional robustness of the hybrid scheme is not required, thus the algebraic scheme is used.

C.   CFD Solver

We use the SUmb flow solver [14]. SUmb is a finite-volume, cell-centered multiblock solver for the compressible Euler, laminar Navier–Stokes, and RANS equations (steady, unsteady, and time-periodic). It provides options for a variety of turbulence models with one, two, or four equations and options for adaptive wall functions. The Jameson–Schmidt–Turkel (JST) scheme [15] augmented with artificial dissipation is used for the spatial discretization. The main flow is solved using an explicit multi-stage Runge–Kutta method along with geometric multi-grid. A segregated Spalart–Allmaras (SA) turbulence equation is iterated with the diagonally dominant alternating direction implicit (DDADI) method. We have developed a discrete adjoint method for the Euler and RANS equations for the efficient computation of gradients required for optimizations [16]. The adjoint implementation supports both the full-turbulence and frozen-turbulence modes, but in the present work we use the full-turbulence adjoint exclusively. The adjoint equations are solved with preconditioned GMRES [17] using PETSc [18, 19, 20]. We have previously performed extensive Euler-based aerodynamic shape [21, 22] and aerostructural optimization [11, 23]. However, we have observed serious issues with the resulting "optimal" Euler-based designs due to the lack of fidelity in the physical model. While Euler-based optimization can provide design insights, we found that the resulting optimal Euler shapes are significantly different from those obtained with RANS [16]. Euler-optimized shapes tend to exhibit non-physical features, such as a sharp pressure recovery near the trailing edge, and thus RANS-based shape optimization is necessary to achieve realistic designs.

D.   Optimization Algorithm

Because of the high computational cost of CFD solutions, it is critical to choose an efficient optimization algorithm that requires a reasonably low number of function calls. Gradient-free methods, such as genetic algorithms, have a higher probability of getting close to the global minimum for cases with multiple local minima. However, slow convergence and the large number of function calls make gradient-free aerodynamic shape optimization infeasible with the current computational resources, especially for large numbers of design variables. Since we require hundreds of design variables, we use a gradient-based optimizer combined with adjoint gradient evaluations to solve the problem efficiently.

The optimization algorithm we use is SNOPT (sparse nonlinear optimizer) [24] through the Python interface py-Opt [25] for all results presented here. SNOPT is a gradient-based optimizer that implements a sequential quadratic programming method; it is capable of solving large-scale nonlinear optimization problems with thousands of constraints and design variables. SNOPT uses a smooth augmented Lagrangian merit function, and the Hessian of the Lagrangian is approximated using a limited-memory quasi-Newton method.

III.   Problem Formulation

The goal of this optimization case is to perform lift-constrained drag minimization of the Common Research Model (CRM) wing using Reynolds-averaged Navier–Stokes equations. In this section, we discuss the problem formulation for this test case.

A.   Initial Geometry

The initial geometry is a wing with a blunt trailing edge extracted from the CRM geometry. The geometry and specifications are given by the Aerodynamic Design Optimization Discussion Group [2]. The fuselage and tail are deleted from the original CRM, and the root of the remaining wing is moved to the symmetry plane. The initial geometry is shown in Figure 2. All coordinates are scaled by the mean aerodynamic chord (275.8 in). The reference point is at $x = 1.2077$ and $z = 0.007669$. The reference area is 3.407014.

Figure 2. Initial geometry scaled by its mean aerodynamic chord.

B.   Grid Convergence Study

We generate the mesh for the BWB using an in-house hyperbolic mesh generator. The mesh is marched out from the surface mesh using an O-grid topology to a farfield located at a distance of 25 times the span. The nominal cruise flow condition is Mach 0.85 with a Reynolds number of 5 million based on mean aerodynamic chord. The grid we generated for the test case optimization contains 28.8 million cells. The grid size and y^+_{max} values at the nominal operating condition are shown in Table 1.

Grid level	Grid size	y^+
L00	$230,686,720$	0.233
L0	$28,835,840$	0.493
L1	$3,604,480$	0.945
L2	$450,560$	2.213

Table 1.   Several grid sizes are used in our optimizations: from 450k to 28.8M cells.


Figure 3.   The grid convergence study shows that the difference between the 28.8M and 230M grids is within 1 drag count.

We perform a grid convergence study to determine the resolution accuracy of this grid. All the grids are generated using an hyperbolic mesh generator with coarse or refined spacing. Figure 3 shows the mesh convergence plot the grid for the initial geometry, as well as the grid for the optimal wing. We can see that the 28.8M grid has sufficient accuracy: The difference in the drag coefficients between the 28.8M and 230M grids is within 1 drag count. The surface symmetry plane for the L0, L1, and L2 grids are shown in Figure 4.

5

(a) L0 grid has 28.8M cells.

(b) L1 grid has 3.6M cells

(c) L2 grid has 450k cells

Figure 4. O-grids of varying sizes were generated using an hyperbolic mesh generator.

6

C. Optimization Problem Formulation

We minimize the drag coefficient by varying the shape design variables subject to a lift constraint ($C_L = 0.5$). In addition, the pitching moment is constrained to be $C_{My} \geq -0.17$. The shape design variables are z-coordinates of 768 control points on the FFD volume, and angle-of-attack. There are 750 thickness constraints imposed in a grid with 25 chordwise and 30 spanwise stations. The thickness is set to by greater than 25% of the initial thickness at each location. Finally, the internal volume is constrained to be greater than or equal to the initial volume.

D. Surface Sensitivity on the Baseline Geometry

To examine the potential improvements of the baseline geometry, we performance a sensitivity analysis of the baseline geometry. We see that the baseline can be improved upon through changes in the sectional airfoil shape. The sensitivity of the drag and pitching moment with respect to the airfoil shape can be visualized through the sensitivity contour plot in Figure 5. Here, we plot the derivatives of C_D and C_{My} with respect to shape variations in the z direction. The regions with the highest gradient of C_D are near the shock on the upper surface. This indicates that shock reduction through local shape changes is the major driver in reducing C_D at the beginning of the optimization. As for C_{My}, the shape changes near the root and tip of the wing are most effective in adjusting pitching moment. However, these sensitivity plots are only a linearization about the current design point, and they provide no information about the constraints. Nonetheless, these sensitivity plots indicate what drives the design at this design point.


Figure 5. Sensitivity study of the baseline shows which shape changes yield the largest improvements.

IV. Aerodynamic Shape Optimization of the CRM Wing

In this section, we present the results of the aerodynamic design optimization of the CRM wing. A grid size of 28.8 million cells is used for the optimization. We use a multilevel optimization acceleration technique to significantly improved the performance of the optimization. The details of this acceleration technique are discussed in Section V.

The optimized wing has 8.5% lower drag compared to that of the baseline wing. The drag decreased from 199.7 counts to 182.8 counts at the nominal flow condition. Figure 6 show the comparison between the baseline and the optimized wing.
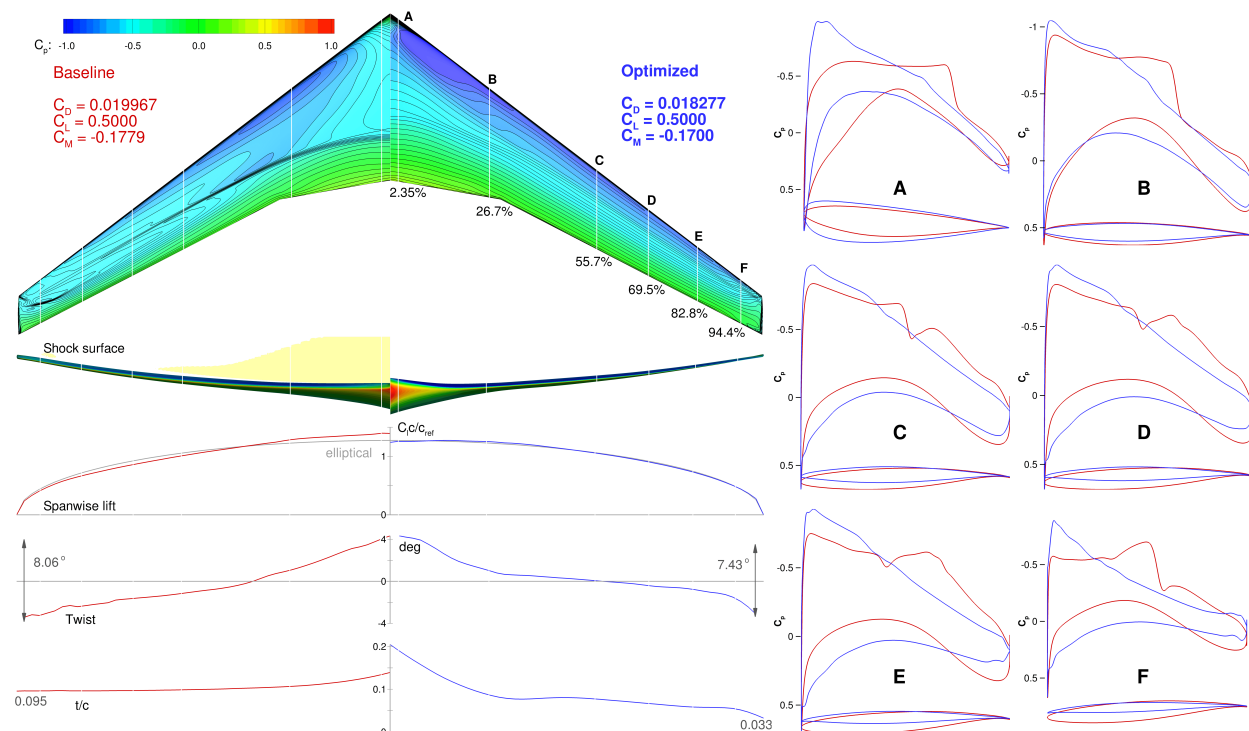

Figure 6. The optimized wing is shock-free and has 8.5% lower drag.

In this figure, the baseline wing results are shown in red and the optimized wing results are shown in blue. At the optimum, the lift coefficient target is met and the pitching moment is reduced to the lowest allowed value. The lift distribution of the optimized wing is much closer to the elliptical distribution, indicating an induced drag that is close to the theoretical minimum. This is achieved by fine-tuning the twist distribution and airfoil shapes. The baseline wing has a near linear twist distribution. The optimized design has more twist at the root and at the tip, and less twist near mid wing. The overall twist angle only changed slightly from 8.06 degrees to 7.43 degrees.

The optimized thickness distribution is significantly different from that of the baseline. Due to the volume constant, the overall volume has to be conserved. Therefore, the optimizer chooses to increase the thickness at the root and decrease the thickness at the tip. The root t/c is over 20%. The low thickness near the tip would in practice incur structural weight penalty. To obtain a more realistic design, we also performed additional optimization with a more strict thickness constraint in Section VI.

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

Another noticeable feature in the optimized wing is the sharp leading edge. The optimizer explores the weakness in the problem formulation. With a single-point optimization, there is no penalty for thinning out the leading edge. However, sharp leading edge airfoils experience adverse performance at off-design conditions, since the flow is prone to separation at off-design angles-of-attack. We explore these issues in more detail and perform a multi-point

8

optimization in Section VII.

V. Multilevel Optimization Acceleration Technique

In this section, we present an acceleration technique to increase the computational efficiency of the aerodynamic shape optimization. Aerodynamic shape optimization is a computational intensive endeavor. The majority of the computational time is spent in the flow solver, and in the gradient calculation. There are three possible improvements can be made to reduce the optimization time.

The first improvement is to reduce the flow solution time. This has be extensively researched by the CFD community. Commonly used methods, such as multigrid, pre-conditioning, and variations on Newton-type methods, can improve the convergence of the solver, thus reducing the overall optimization time.

The second improvement is to reduce the gradient computational time, which was pioneered by Jameson [1] through the development of adjoint methods, which efficiently compute gradients with respect to large numbers of shape design variables. With an efficient adjoint implementation, the cost of computing the gradient of a single function of interest with respect to hundreds or thousands of shape design variables is roughly the cost of one flow solution [16]. For a generalization of the adjoint method and its connection to other sensitivity analysis methods, see [26].

The third improvement that can be made is to reduce the number of function and gradient calls. This can be achieved by using a surrogate model [23]. However, as the dimension of the problem increases, surrogate modeling becomes less effective and less accurate.

In this paper, we present a method that is inspired by the multigrid method in CFD. We use smaller grids to accelerate the convergence of a large grid. Since it is less costly to compute both flow solution and the gradient in a smaller grid, we perform the optimization first on a smaller first until a certain level of optimality is achieved. Then, we move on to the next grid level and use the design variables from the previous grid level as the initial design variables. Since the drag and lift coefficients are generally different on each grid level, the approximate Hessian (used by the gradient-based optimizer) must be restarted. This process is repeated until the last grid level has converged.

We demonstrate this method using the CRM wing optimization case solved in the previous section. Three grid levels are used: L2 (451k), L1 (3.6M) and L0 (28.8M). The merit function, optimality, and feasibility history is shown in Figure 7. We can see that the majority of the iterations are performed on the coarse grid and as a result, the number of the function and gradient evaluations on the successively finer grids is greatly reduced. Table 2 summarizes the computational time spent on each grid level. Thanks to the optimization with the coarser grids, only 18 iterations are needed on the L0 grid. However, L0 grid still uses the most of the computational cost in terms of proc-hr. Due to the high cost of flow and adjoint solution on L0 grid, as well as limited computational resources, we can not perform an optimization with only L0 grid. Assuming the same number of iterations used for the L2 grid (638) would be needed for the L0 grid, the computational cost would be 23 times higher than the multilevel approach, which would correspond to 16 days using 1248 processors.

Grid level	Iterations	Procs	Time (hr)	Total proc-hr
L2	638	64	29.3	1875.2
L1	89	256	20.2	5171.2
L0	18	1248	11.1	$13,852.8$

Table 2. The number of iterations on the L0 grid is reduced to 18.

Figure 8 shows the initial and optimized results at each grid level. If we examine the results more closely, we see that the optimized results of L2, L1, and L0 are all similar to each other. This validates the underlying assumption of this method: that a coarser grid provides a good approximation to the design space of the finer grid. The majority of the computational efforts on the subsequent grid levels are spent on smoothing out the shock that re-appeared due to the finer grid spacing. This multilevel acceleration technique proved to significantly reduce the number of iterations needed to optimize in the fine grid, and the total computational effort was greatly reduced.

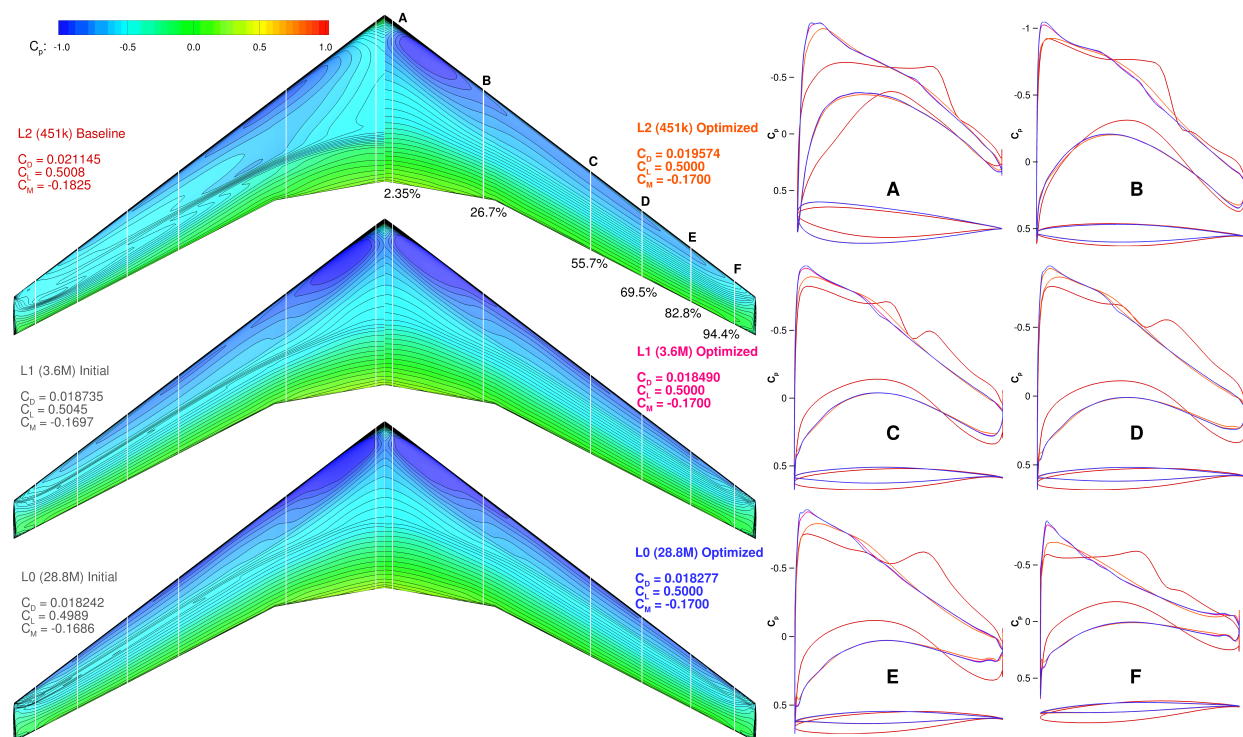**Figure 7. The majority computations are performed on the coarse grid.**

Figure 8. The optimized results of each grid level exhibit only subtle differences.

VI.   CRM Wing Optimization with 100% Thickness Constraint

In Section IV, the optimized wing has a thickened root airfoil and an unrealistically thin tip airfoil. To address this issue, we performed an optimization with the same setup except for modified thickness constraints: All airfoil thickness constraints must be greater than or equal to the initial thickness. The optimization is performed on the L2 grid, and the results are shown in Figure 9.
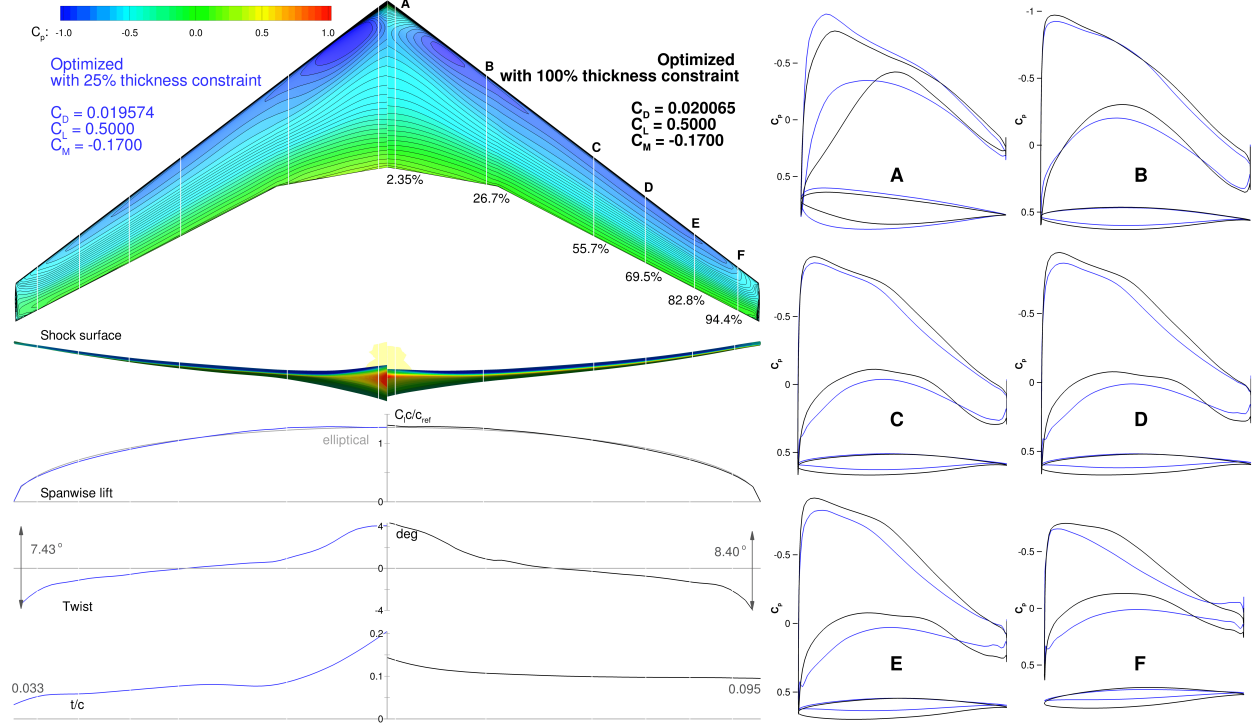

Figure 9. The drag on the optimized wing is 4 count higher if no airfoil thickness reduction is allowed.

The results of the optimization with 100% thickness constraint are shown in black. The spanwise lift and twist distributions for both cases are similar to each other. However, the pressure distribution and airfoil shapes are significantly different, especially those near wing root and wing tip. As a result, the optimized wing with 100% thickness constraint has 5 additional drag counts when compared the wing optimized subject to the 25% thickness constraint. Such aerodynamic performance penalty may be compensated by the reduction in wing structural weight. A detailed aerostructural optimization would be necessary to examine the tradeoffs, and such an optimization has been performed for a similar wing by Kennedy et al. [27].

VII.   Multi-point Aerodynamic Shape Optimization of the CRM Wing

Transport aircraft operate at multiple cruise conditions due to variability in the missions and air traffic control restrictions. Single-point optimization at the nominal cruise condition could inflate the benefit of the optimization, since it is likely to improve the on-design performance while reducing the performance under off-design conditions. In Section IV, the single-point optimized wing exhibited an unrealistically sharp leading edge near the wing tip. This was caused by a combination of the volume constraint and the single-point formulation. A sharp leading edge is prone to flow separation at off-design conditions. We attempt to address this issue by performing a multi-point optimization. The optimization is performed on the L2 grid. We choose five equally weighted flow conditions by varying the lift coefficient and the Mach number. The flight conditions are the nominal cruise, $\pm 10\%$ of cruise C_L, and ± 0.01 of cruise Mach, as shown in Table 3. More sophisticated ways of choosing multi-point flight conditions can be used,

such as an automated selection of the points that minimize fleet-level fuel burn [23].The objective function is the average drag coefficient of five flow cases.

Flow Case	C_L	Mach number
1	0.50	0.85
2	0.55	0.85
3	0.45	0.85
4	0.50	0.84
5	0.50	0.86

Table 3. The flow cases are arranged in a five-point stencil in Mach-C_L space.
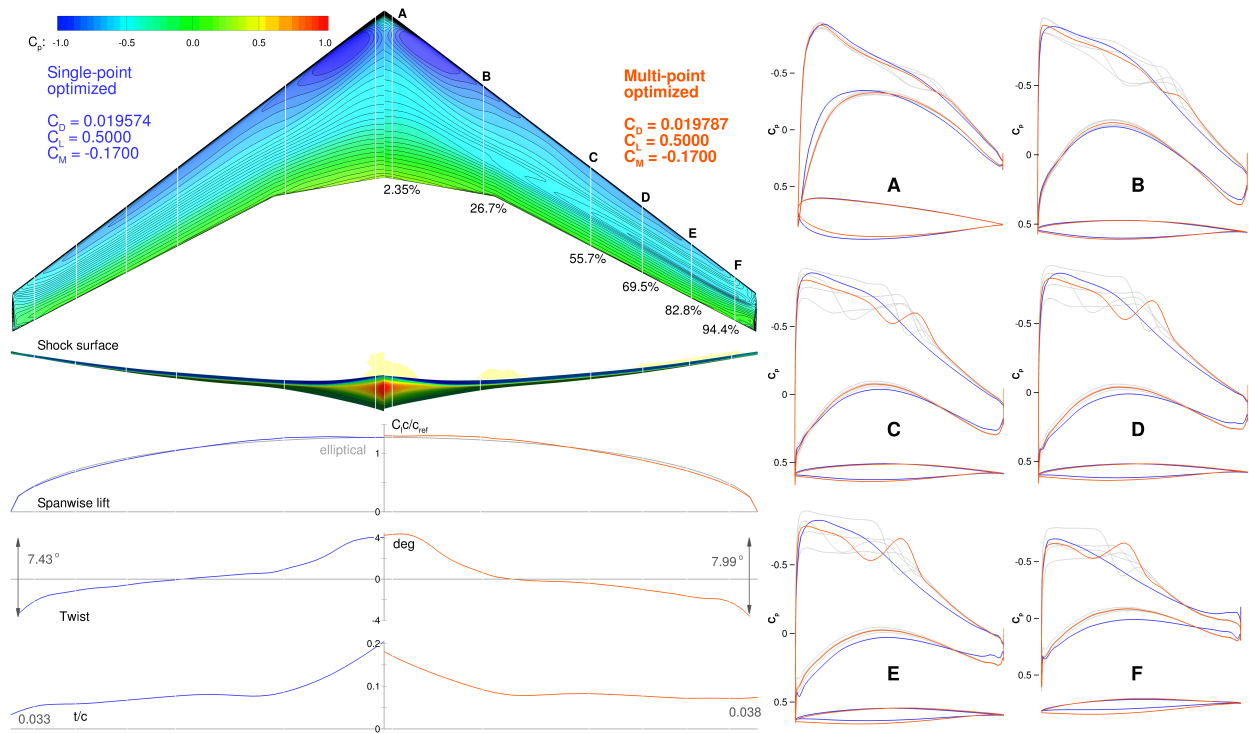
Figure 10. The multi-point optimized wing has a localized weak shock on the upper surface.

The comparison between the single-point and multi-point optimized designs is shown in Figure 10. The single-point results are shown in blue and the multi-point results are shown in orange. Note that only the results at the nominal flight condition are plotted. In the multi-point optimization, the sectional C_p of flow case 2-5 are plotted in gray. Compared to the single-point optimization, the multi-point optimized wing has a localized weak shock at all flow conditions. The leading edge is less sharp than that of the single-point optimized wing. Additional flow cases, such as a low-speed flight condition, would be needed to further improve the leading edge. The overall pressure distribution of the multi-point design is similar to that of the single-point design. The twist and lift distributions are nearly identical. Most of the differences are in the chordwise C_p distributions in the outer wing section. The drag coefficient at nominal condition is approximately 2 count higher. However, the performance at the off-design conditions is significantly improved.

To visualize this improvement, and to better understand the effects of multi-point optimization, we plot ML/D contours of the baseline, single-point, and multi-point designs with respect to C_L and cruise Mach in Figure 11 .

ML/D provides a metric for quantifying aircraft range based on the Breguet range equation with constant thrust specific fuel consumption. While the thrust specific fuel consumption is actually not constant, assuming it to be constant is acceptable when comparing range performance in a limited Mach number range [28]. We added 100 drag counts to the computed drag to account for the drag due to the fuselage, tail and nacelles and get more realistic ML/D values.
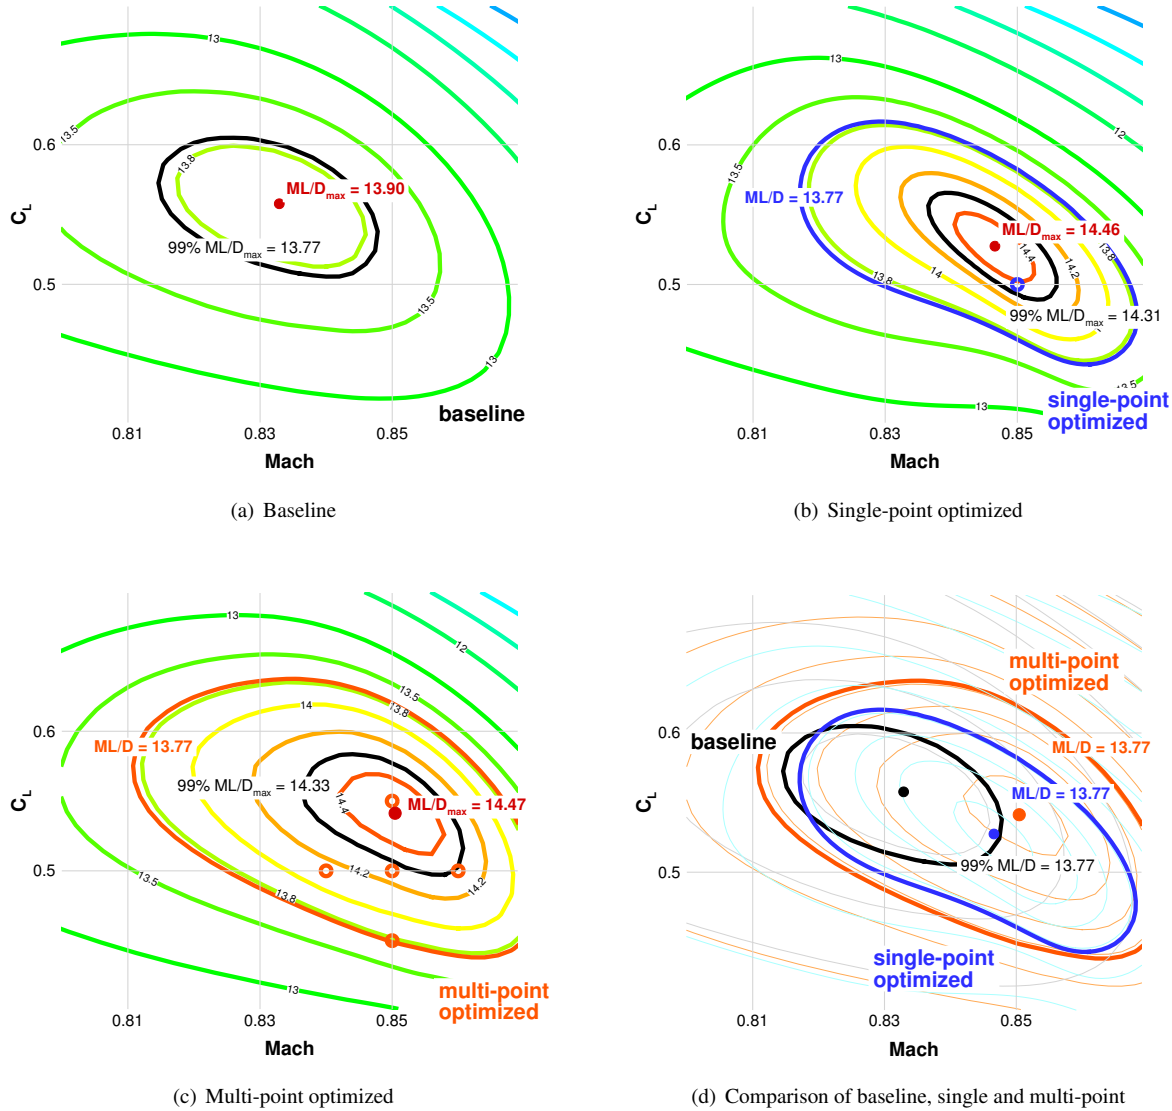

(a) Baseline

(b) Single-point optimized

(c) Multi-point optimized

(d) Comparison of baseline, single and multi-point

Figure 11. The multi-point optimized wing has better off-design performance and is more robust.

The baseline maximum ML/D is at a lower Mach number and a higher C_L compared to the nominal flight condition. The single-point optimization significantly increases the maximum ML/D and the ML/D at the operation condition. In addition, the maximum ML/D occurs much closer to the nominal cruise condition. The shapes of the contours are also altered to move the maximum toward the cruise flight condition. For fixed $C_L = 0.5$, the maximum ML/D occurs at the nominal Mach of 0.85, which is equivalent to the drag bucket in a drag divergence plot. For the multi-point optimization, the flight conditions for optimization are spread in the Mach-C_L space, resulting in a flat-

tened ML/D variation near the maximum, resulting in more uniform performance for a range of flight conditions. The 99% ML/D contour is also larger than that of the single-point optimum. By performing a multi-point optimization, we achieve a more robust design and increase the overall performance at both on- and off-design conditions.

VIII.   Investigation on the Number of Shape Design Variables

With an efficient adjoint implementation, the cost of computing gradients is nearly independent of the number of design variables. Therefore, we took advantage of this efficiency by using a larger number of design variables in the optimizations presented so far. Now, we examine the effect of reducing the number of design variables on the optimized wing. We perform the same optimization as in Section IV with on fourth of the design variables: 192 as opposed to the original 768. This reduction is achieved by halving the number of control point in both spanwise and chordwise directions. The L2 grid is used for this study, and the comparison with the higher-dimensionality case is shown in Figure 12.
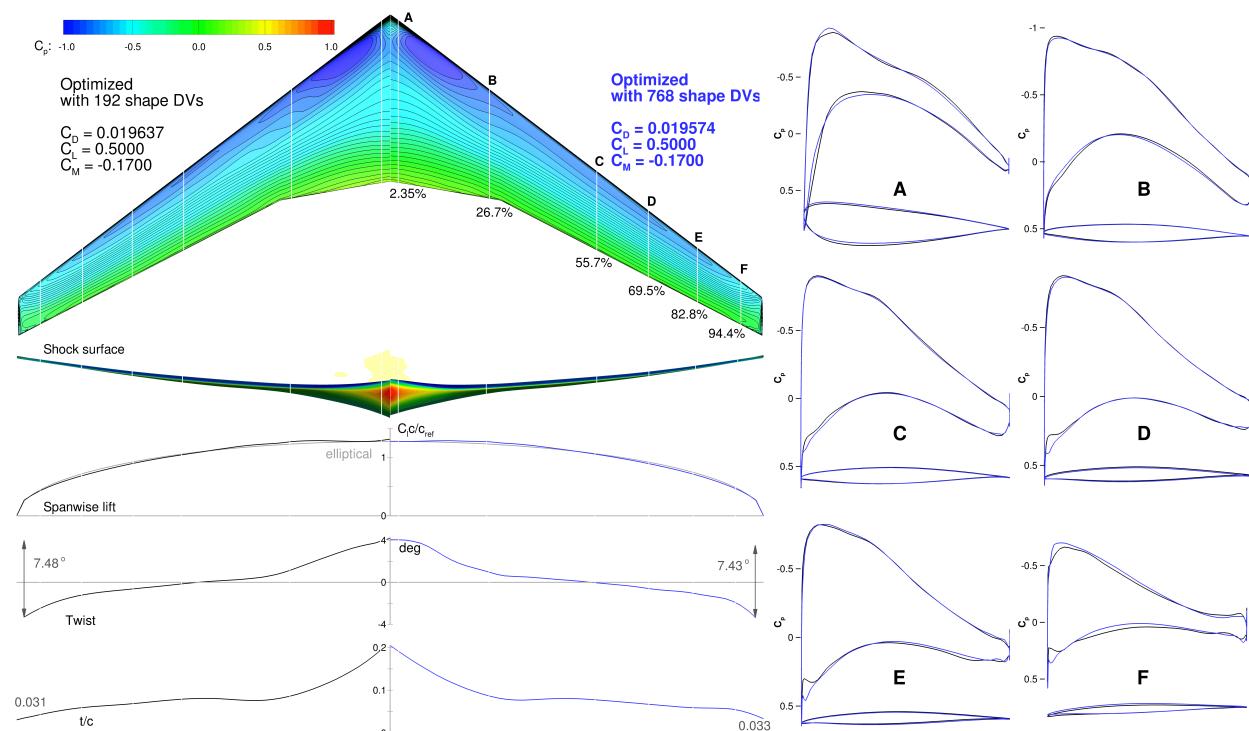

Figure 12. The optimized wing with 192 design variables is similar to that with 768 design variables.

We see that both optimization results have a similar near-elliptical lift distribution, twist, and thickness distributions. The sectional airfoil shapes are also similar. With a higher number of design variables, the optimizer can has a finer control over the geometry. The effects can be seen on the sectional C_p on sections D, E, and F near the leading edge. The difference in drag coefficient is about 0.6 count. Therefore, we conclude that an adequate optimized design can be achieved with a smaller number of design variables. For an optimization process whose computational cost that scales with the number of design variables, such as gradients computed with finite-differences, or when using gradient-free optimizers, using a smaller number of design variables would be beneficial and would not have a significant impact on the optimized design.

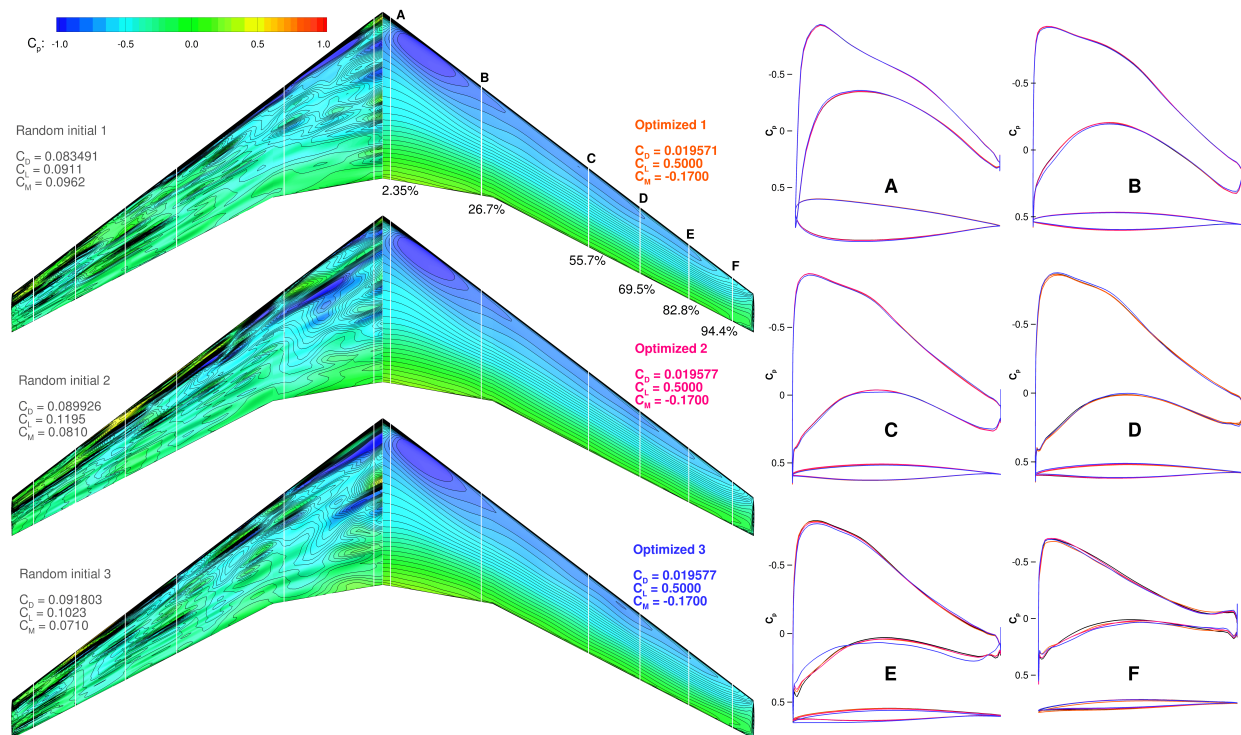**Figure 13. The initial geometries are randomly generated based on the baseline CRM wing.**

Figure 14. All three optimizations with random starting geometries converged to similar optima.


Figure 15. The merit function values between optimized designs show the local minima.


Figure 16. The distances between the optimized wings in the design space are similar.

IX.   Investigation on the Multi-modality of Aerodynamic Shape Optimization Using Random Initial Geometries

High-fidelity aerodynamic shape optimization with a large number of design variables has the potential to have multiple local minima. The problem is that due to the high number of dimensions, the design space is difficult to visualize and it is computationally costly to find multiple local minima. We explore the multi-modality of this aerodynamic shape optimization problem by performing four separate optimization, where each optimization starts with a different geometry. The first run is the nominal optimization shown in Section IV. The other three runs has the same optimization formulation with randomly generated initial geometries. The random geometries are generated based on the CRM wing with a random surface perturbation for each shape design variables, resulting in completely different geometries. The initial starting points for the three random runs are shown in Figure 13.

Figure 17 shows the optimized results from a random initial starting geometry. The optimization is performed on the L2 grid. We can see that the performance of the initial design is extremely poor. This is no surprise, since the airfoil shapes are unlike anything one would design: they exhibit oscillations and sharp edges, resulting in wildly varying C_p distribution. In spite of these wild shapes, the optimizer is able to smooth out the airfoils and achieve a shock-free wing similar to the original single-point design presented in Section IV.


Figure 17. The optimization manages to start from a random geometry and converge to an optimal wing that is shock free.

We performed the same optimization for three random starting points and compared the results against each other, as well as against the single-point optimized wing, as shown in Figure 14. Each random optimized result is color-coded, and the nominal optimized result from Section IV is shown in black. Overall, there are only small differences between the four design as evidenced by the similar C_p distributions and cross sectional shapes. The difference in drag between all four designs is within a drag count. However, there are still some small visible differences, indicating the possibility of local minima.

To further visualize this design space, we compute the merit function in the design space between two optimized designs, as shown in Figure 16. The merit function is a combination of the objective function and the constrains. We are able to visualize a slice of the design space by plotting the merit function along a line between two optima. As

17

shown in this figure, the merit function does appear to have local minima, even though the values of the merit function are within one count among those optima. We believe that the design space of this aerodynamic shape optimization problem is relatively flat with many humps and local minima within this flat region. The humps and local minima could also be caused by the constraints.

X. Conclusions

The optimization results of the CRM wing are presented for the Aerodynamic Design Optimization Workshop. The drag coefficient is minimized subject to lift, pitching moment, and geometric constraints. The optimization is performed on a grid with 28.8 million cells using a total of 768 shape design variables. The drag coefficient of the optimized design is reduced by 8.5%, from 199.7 counts to 182.8 counts. The optimization used a multilevel acceleration technique that significantly reduced the total computational time.

We found that the optimized design exhibits a small thickness-to-chord ratio (3.3%) at the tip, which would incur a large structural weight penalty in a real wing. Thus, additional optimization was performed with 100% thickness constraints, resulting in a increase of drag by 5 counts for the optimized design. However, the associated reduction in structural weight would likely off-set this penalty in a real wing.

We also presented a multi-point optimization of the CRM wing. This resulted in a more robust design than that of the single-point optimization, as evidenced by the enlarged contour of the 99% maximum ML/D. We also compared the contours of ML/D for the baseline, single point optimum and multi-point optimum. Both single-point and multi-point optimization shifted the maximum ML/D toward the nominal flight condition. A more robust design and improved the overall performance at both on- and off-design conditions were achieve using multi-point optimization.

The multi-modality of the aerodynamic shape optimization problem was examined by optimizing randomly generated initial geometry. The optimized wings had similar airfoil shapes and the drag coefficients were converged within 1 count of difference. However, we observed subtle differences in the airfoil shapes. The merit function values between the optimized wings were plotted to visualize the design space. We conclude that the design space is relatively flat and contains multiple local minima.

XI. Acknowledgments

References

[1] Jameson, A., "Aerodynamic Design via Control Theory," Vol. 3, No. 3, 1988, pp. 233–260. doi:10.1007/BF01061285.

[2] Osusky, L. M., *A Novel Numerical Tool for Aerodynamic Shape Optimization in Turbulent Flow*, Ph.D. thesis, University of Toronto, 2013.

[3] Reuther, J. J., Jameson, A., Alonso, J. J., Rimlinger, M. J., and Saunders, D., "Constrained Multipoint Aerodynamic Shape Optimization Using an Adjoint Formulation and Parallel Computers, Part 1," *Journal of Aircraft*, Vol. 36, No. 1, 1999, pp. 51–60. doi:10.2514/2.2413.

[4] Reuther, J. J., Jameson, A., Alonso, J. J., Rimlinger, M. J., and Saunders, D., "Constrained Multipoint Aerodynamic Shape Optimization Using an Adjoint Formulation and Parallel Computers, Part 2," *Journal of Aircraft*, Vol. 36, No. 1, 1999, pp. 61–74. doi:10.2514/2.2414.

[5] Nemec, M., Zingg, D. W., and Pulliam, T. H., "Multipoint and Multi-Objective Aerodynamic Shape Optimization," *AIAA journal*, Vol. 42, No. 6, 2004, pp. 1057–1065. doi:10.2514/1.10415.

[6] Lyu, Z. and Martins, J. R. R. A., "RANS-based Aerodynamic Shape Optimization of a Blended-Wing-Body Aircraft," *43rd AIAA Fluid Dynamics Conference and Exhibit*, June 2013.

[7] Lyu, Z. and Martins, J. R. R. A., "Aerodynamic Shape Optimization Studies of a Blended-Wing-Body Aircraft," *Journal of Aircraft*, 2014, (Accepted).

[8] Holst, T. L. and Pulliam, T. H., "Aerodynamic Shape Optimization using a Real-Number-Encoded Genetic Algorithm," *NASA Technical Reports*, 2001.

[9] Sasaki, D., Morikawa, M., Obayashi, S., and Nakahashi, K., "Aerodynamic Shape Optimization of Supersonic Wings by Adaptive Range Multiobjective Genetic Algorithms," 2001, pp. 639–652. doi:10.1007/3-540-44719-9_45.

[10] Chernukhin, O. and Zingg, D. W., "Multimodality and Global Optimization in Aerodynamic Design," *AIAA Journal*, April 2013, pp. 1–13. doi:10.2514/1.J051835.

[11] Kenway, G. K. W., Kennedy, G. J., and Martins, J. R. R. A., "A Scalable Parallel Approach for High-Fidelity Steady-State Aeroelastic Analysis and Adjoint Derivative Computations," *AIAA Journal*, 2013, (In Press).

[12] Kenway, G. K. W. and Martins, J. R. R. A., "Multi-point High-fidelity Aerostructural Optimization of a Transport Aircraft Configuration," *Journal of Aircraft*, 2013, (In Press).

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

[14] van der Weide, E., Kalitzin, G., Schluter, J., and Alonso, J., "Unsteady Turbomachinery Computations Using Massively Parallel Platforms," *44th AIAA Aerospace Sciences Meeting and Exhibit*, 2006. doi:10.2514/6.2006-421.

[15] Jameson, A., Schmidt, W., and Turkel, E., "Numerical Solution of the Euler equations by Finite Volume Methods Using Runge Kutta Time Stepping Schemes," *14th AIAA, Fluid and Plasma Dynamics Conference*, 1981.

[16] Lyu, Z., Kenway, G., Paige, C., and Martins, J. R. R. A., "Automatic Differentiation Adjoint of the Reynolds-Averaged Navier-Stokes Equations with a Turbulence Model," *43rd AIAA Fluid Dynamics Conference and Exhibit*, June 2013.

[17] Saad, Y. and Schultz, M. H., "GMRES: A Generalized Minimal Residual Algorithm for Solving Nonsymmetric Linear Systems," *SIAM Journal on Scientific and Statistical Computing*, Vol. 7, No. 3, 1986, pp. 856–869. doi:10.1137/0907058.

[18] Balay, S., Gropp, W. D., McInnes, L. C., and Smith, B. F., "Efficient Management of Parallelism in Object Oriented Numerical Software Libraries," *Modern Software Tools in Scientific Computing*, edited by E. Arge, A. M. Bruaset, and H. P. Langtangen, Birkhäuser Press, 1997, pp. 163–202. doi:10.1007/978-1-4612-1986-6_8.

[19] Balay, S., Brown, J., , Buschelman, K., Eijkhout, V., Gropp, W. D., Kaushik, D., Knepley, M. G., McInnes, L. C., Smith, B. F., and Zhang, H., "PETSc Users Manual," Tech. Rep. ANL-95/11 - Revision 3.4, Argonne National Laboratory, 2013.

[20] Balay, S., Brown, J., Buschelman, K., Gropp, W. D., Kaushik, D., Knepley, M. G., McInnes, L. C., Smith, B. F., and Zhang, H., "PETSc Web page," 2013, http://www.mcs.anl.gov/petsc.

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

[22] Lyu, Z. and Martins, J. R. R. A., "Aerodynamic Shape Optimization of a Blended-Wing-Body Aircraft," *51st AIAA Aerospace Sciences Meeting including the New Horizons Forum and Aerospace Exposition*, 2013. doi:10.2514/6.2013-283.

[23] Liem, R. P., Kenway, G. K., and Martins, J. R. R. A., "Multi-Point, Multi-Mission, High-Fidelity Aerostructural Optimization of a Long-Range Aircraft Configuration," *Proceedings of the 14th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*, Indianapolis, IN, Sep 2012. doi:10.2514/6.2012-5706.

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

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

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

[27] Kennedy, G. J., Kenway, G. K., and Martins, J. R. R. A., "Optimal Aerostructural Tradeoffs for Advanced Composite Wings," *AIAA Science and Technology Forum and Exposition (SciTech), National Harbor, MD*, 2014.

[28] Torenbeek, E., *Advanced Aircraft Design: Conceptual Design, Technology and Optimization of Subsonic Civil Airplanes*, Aerospace Series, Wiley, 2013. doi:10.1002/9781118568101.