

High-Fidelity Wing Design Exploration with Gradient-Based Optimization


by

Nicolas Peter Bons

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Aerospace Engineering)
in The University of Michigan
2020

Doctoral Committee:

Professor Joaquim R. R. A. Martins, Co-Chair
Dr. Charles A. Mader, Co-Chair
Professor Carlos E. S. Cesnik
Professor Yin Lu Young



Nicolas Peter Bons
nbons@umich.edu
ORCID iD: 0000-0002-4743-8726

© Nicolas Peter Bons 2020

This dissertation is dedicated to Brooke, Rose, and Lily, for always
making the homecoming the best part of the day.

Acknowledgments

Early on in graduate school, it struck me that in an age of instant virtual communication, the demand for air travel should go down, which would put me in a twilight industry. We know from experience though, that the opposite is true (neglecting the COVID-19 anomaly). I came to the conclusion that we humans have an innate need for physical connection that cannot be satisfied in the virtual world. Furthermore, I decided that researching ways to improve aircraft was a worthwhile investment of my time and effort if it contributed to helping people have those face-to-face interactions. So, at the beginning of this dissertation, I want to acknowledge those people who have helped me along the way and filled my life with connection.

First, I want to thank Prof. Martins for welcoming me into his research group and having faith in my developing ability. I had the unique opportunity of working closely with Prof. Martins on two trips to Brazil and while doing consulting work in Palo Alto—experiences I won't soon forget. I continue to be impressed with his enthusiasm for self-improvement and the effort he puts into building relationships.

There are numerous people in the MDO Lab who helped me get to where I am today. All of the work in this dissertation is built on the foundation laid by those who came before me, and I am in awe of all they accomplished. I especially want to thank Sandy Mader for being a patient, nonjudgmental teacher and for his calm, methodical approach to troubleshooting. I was very fortunate to start my graduate studies with three outstanding guys. I want to thank John for being my sounding board for nearly five years, but more than that for being an inspiring example of selflessness and kindness. I greatly appreciated Sham's shared interest in literature and his compulsion to understand things deeply. Last but not least, I want to thank Justin for his good humor, expansive knowledge, and for always checking in to see how the family is doing. I owe a lot to my office mates—John, Ney, Tim, John H., Mohamed, Gurmeet, Xiaosong, Ali, Josh, Ben, and Anil—for putting up with me and for making the office a pleasant place to be. I also want to thank Eirikur and Neil for liberally sharing their technical expertise with me. To everyone who is or has been a part of the MDO Lab, thank you for contributing to the genuine camaraderie that exists in the group. I think that it is a significant factor in our

continuing success.

I am also grateful for friends and family who buoyed me up and supported my family during the past five years—we couldn't have made it without you. To Mom and Dad, thank you for always believing in me and for patiently encouraging me to reach my highest potential. Finally, the greatest debt of gratitude is owed to my wife, Brooke. Thank you for your compassion and optimism and for patiently enduring the difficult parts of the PhD years. I look forward to a bright future with you and our children.

Table of Contents

Dedication	ii
Acknowledgments	iii
List of Figures	vii
List of Tables	xi
List of Symbols	xii
List of Acronyms	xiv
Abstract	xv
Chapter 1 Introduction	1
1.1 A Brief History of Wing Design	2
1.2 Numerical Optimization of Wings	4
1.3 Motivation and Objectives	7
1.4 Contributions and Outline	9
Chapter 2 Overview of the MACH Framework	10
2.1 Optimization Algorithm	12
2.2 Geometry	12
2.3 High-fidelity Physics Modeling	14
2.4 Mission Analysis	16
Chapter 3 Exploration of Aerodynamic Wing Design	20
3.1 Methodology	21
3.2 Parametrization Studies	25
3.3 Full Case Optimization	40
3.4 Sweep variable study	43
3.5 Summary	44

Chapter 4	Exploration of Aerostructural Wing Design	48
4.1	Problem Description	48
4.2	Methods	50
4.3	General Optimization Problem	52
4.4	Aerodynamic Shape Optimization	56
4.5	Multi-level Optimization Procedure	60
4.6	Single-point Optimization	61
4.7	Robust Design Optimization	66
4.8	Summary	72
Chapter 5	Practical Wing Design Optimization	74
5.1	Problem Description	76
5.2	Computational Framework	80
5.3	Simple Wing Optimization	83
5.4	Full Aircraft Optimization	89
5.5	Summary	94
Chapter 6	Final Remarks	96
6.1	Conclusions	96
6.2	Contributions	98
6.3	Recommendations for Future Work	100
Bibliography		102

List of Figures

Figure

2.1	XDSM diagram of aerostructural optimization with MACH.	11
2.2	The geometric parametrization uses FFD control points for local control and a B-spline reference axis for global control.	13
2.3	Various point-based geometric constraints are possible in the MACH framework. Linear constraints can also be added to the FFD control points.	14
2.4	Representative depiction of smeared stiffness model.	15
3.1	Grid convergence study at $M = 0.5$, $Re = 5 \times 10^6$ and $C_L = 0.2625$	22
3.2	Geometric parametrization of the aerodynamic multimodality optimization problem. The wing surface is embedded in the FFD volume. Design variable definitions are indicated with arrows.	23
3.3	Elliptical distribution can be reached with different numbers of FFD control sections and distributions.	26
3.4	Wings optimized with respect to twist variables generate an elliptical lift distribution.	26
3.5	Elliptical planform and lift distributions obtained from optimization with respect to chord variables.	28
3.6	The viscous drag on a flat plate in laminar flow can be reduced by increasing the average BL thickness over the plate. ($Re = 10^6$)	29
3.7	The optimal chord distribution for viscous flow is heavily dependent on C_L	30
3.8	Optimization using OpenAeroStruct confirms multiple optimal chord distributions at low C_L	31
3.9	(a) Three optimizations from random chord distributions yield two local minima. (b) Varying the number of chord variables changes the optimal distribution slightly.	32
3.10	Forcing the chord distribution to decrease monotonically increases drag by less than one count.	32
3.11	Optimizing from 10 unique starting points with respect to chord and twist variables using Euler analysis yields a single optimum.	33

3.12	An optimization with both twist and chord variables tolerates more variation in the chord distribution at a lower C_L . As C_L increases, the planform shape oscillations disappear.	34
3.13	Comparison of different quantities and spanwise distributions of control points for an optimization with dihedral variables.	35
3.14	Results of optimizing the dihedral distribution first with L3 meshes and then with L2 meshes.	36
3.15	We use OpenAeroStruct to quickly explore the design space of the dihedral and twist optimization.	38
3.16	(a) A nonplanar wing forms as C_L increases. (b) The chord varies at low C_L and the winglet forms at a high C_L	39
3.17	The optimizer converges on two local minima for $C_L = 0.5$	39
3.18	RANS optimization results starting from the baseline rectangular wing at different C_L values.	40
3.19	Three local minima found with all variables active.	41
3.20	Two local minima found when monotonically decreasing chord is enforced.	43
3.21	One optimal design found when chord and sweep distributions are constrained to be linear.	44
3.22	Comparison of scimitar and straight wing geometries and C_p distributions.	45
3.23	Comparison of spanwise variation on scimitar and straight wings.	45
4.1	Rectangular wing definition	49
4.2	CFD meshes of baseline wing.	51
4.3	The structural mesh has 16,672 2nd-order shell elements.	51
4.4	Geometric parametrization for rectangular wing optimization problem. On the left, the wing surface embedded in the FFD volume. On the right, the CRM planform is reproduced by modifying the FFD control points.	52
4.5	Wave drag can be reduced by sweeping the wing forward or backward, resulting in a multimodal problem. Each of the plotted points is the result of an aerodynamic shape optimization at a fixed sweep angle.	58
4.6	Planforms	59
4.7	Comparison of three aerodynamic shape optimizations.	60
4.8	For this problem, the multi-level approach achieves basically the same design as the single-level optimization on the finest grid—at 60% of the computational cost.	61
4.9	Each successive level of mesh refinement yields additional design changes.	62
4.10	Drag convergence study for the baseline and single-point optimized wings. For the baseline wing, the drag increases as the mesh is refined because the shock is resolved more accurately on finer grids. The improvements on the optimized wings are preserved as the mesh is refined.	63

4.11	Comparison of single-point optimized designs. The optimizations starting from the plank and CRM planforms converge to nearly the same design.	64
4.12	The difference in initial and final design variables for wing optimization problems starting from a plank planform and a Common Research Model (CRM) planform.	65
4.13	Analysis point for low-speed, high-lift separation constraint is placed at the boundary of the climb profile for the Boeing 777-200ER.	68
4.14	Optimizing for robust performance in cruise <i>and</i> climb incurs ~1000 kg increase in cruise fuel burn.	69
4.15	The multipoint design is very similar to the single-point design, but including the low-speed separation constraint elicits striking modifications.	69
4.16	The low-speed separation constraint improves robustness at both cruise and climb flight conditions.	71
4.17	Separation on the upper surface at 9 degrees angle of attack is nearly eliminated with the low-speed separation constraint.	72
5.1	Computational grids for the R-jet configuration.	77
5.2	Computational grids for the R-wing configuration juxtaposed with the R-jet grids.	78
5.3	Basic mission profile.	79
5.4	The ratio of climb fuel burn to cruise fuel burn grows exponentially as mission range decreases. The left column shows data for five different payload weights and the right column shows data for four different cruise altitudes. The data corresponding to the R-jet aircraft description is highlighted in blue.	80
5.5	Points are extracted from the R-wing and R-jet wings in order to create a basic representation of the geometry in pyConcept.	83
5.6	On both baseline and optimized designs, the low-fidelity drag polar from pyConcept coincides with the high-fidelity aerostructural result at cruise flow conditions. The error in the low-fidelity method grows as the speed and altitude decrease.	84
5.7	Pareto fronts from optimizations using different methods of fuel burn calculation.	87
5.8	The low-fidelity aerodynamic analysis used in MissionAnalysis underestimates the drag reduction at several points along the climb and descent profiles.	88
5.9	Comparison of airfoils and C_p distributions for different cases.	92
5.10	Comparison of spanwise lift, twist, and t/c distributions for different cases.	93

5.11 Views of planform and deflected wing shape for the optimized wings. The S+T wing planform is identical to the baseline wing.	94
5.12 Comparison of baseline and optimized wingbox structures. The wing deflection is normalized by the maximum tip deflection on the baseline wing.	95

List of Tables

Table

3.1	Baseline geometry performance at $C_L = 0.2625$	21
3.2	ADODG Case 6 Optimization Problem Statement	25
3.3	Twist optimization results	26
3.4	Chord optimization results	28
3.5	Chord and twist optimization results	34
3.6	Dihedral and twist optimization results. The $C_L = 0.8$ result was obtained by regenerating the volume mesh to fix problems with the formation of negative volumes.	36
3.7	Full optimization problem results. Percentage change is referenced to the baseline wing.	42
3.8	Sweep optimization problem results. Percentage change is referenced to the baseline wing.	44
4.1	Rectangular wing specifications	49
4.2	Aircraft specifications	49
4.3	Grid dimensions	50
4.4	Flow conditions for high-fidelity analyses	52
4.5	Rectangular wing aerostructural optimization problem description.	53
4.6	Aerodynamic shape optimization problem description.	57
4.7	Results of single-point optimization.	62
4.8	Optimization results.	73
5.1	R-jet aircraft Specifications	76
5.2	The R-wing mesh yields a 1500x decrease in proc-hours for a single aerostructural analysis when compared with the R-jet overset mesh.	77
5.3	Mission details	79
5.4	Four methods for calculating the fuel burn and range of a mission.	81
5.5	Description of the full R-wing optimization problem. The quantities for each design variable correspond to the green curve in Figure 5.7.	85
5.6	R-jet optimization problem	91
5.7	Relative difference between baseline and optimized R-jet designs.	92

List of Symbols

α	Angle of attack
λ_{stiff}	Stiffener height
ρ_{∞}	Freestream density
b	Wingspan
C_D	Drag coefficient
C_L	Lift coefficient
C_M	Moment coefficient
c_T	thrust-specific fuel consumption (TSFC)
h_{stiff}	Stiffener height
t_{panel}	Panel thickness
t_{stiff}	Stiffener thickness
V_{∞}	Freestream velocity
\mathcal{R}	Aspect ratio

List of Acronyms

- AD automatic differentiation 14
- ADODG Aerodynamic Design Optimization Discussion Group 5, 20, 24, 39, 40
- AIAA American Institute of Aeronautics and Astronautics 5, 20
- ANK approximate Newton–Krylov 14
- ASO aerodynamic shape optimization 5, 6, 9, 13, 56, 66, 99
- CAD computer-aided design 5, 76
- CAE computer-aided engineering 5
- CAS calibrated airspeed 81
- CFD computational fluid dynamics 2, 4–6, 14, 48, 50, 53, 54, 57, 64, 76, 81, 82, 97
- CRM Common Research Model 20, 48, 50, 53, 56, 58, 63, 67, 68, 74, 75
- CSM computational structural mechanics 6
- DOC direct operating cost 16
- DPW Drag Prediction Workshop 53
- FEA finite element analysis 6
- FFD free-form deformation 12–14, 18, 21, 22, 24, 27, 33, 34, 50, 54, 55, 58, 61, 63, 82, 86, 90
- FSI fluid-structure interaction 6
- KS Kreisselmeier–Steinhauser 15, 55, 100
- MAC mean aerodynamic chord 50, 82
- MACH MDO of aircraft configurations with high fidelity 9, 10, 13–18, 48, 49, 74, 79, 81, 94, 98

MDAO multidisciplinary analysis and optimization 6, 74

MDO multidisciplinary design optimization 6–9, 20, 74, 88, 89, 96, 97, 100, 101

MITC mixed interpolation of tensorial components 76

NACA National Advisory Committee for Aeronautics 2, 3, 21

NASA National Aeronautics and Space Administration 4

NK Newton–Krylov 14

OEM original equipment manufacturer 16

OEW operating empty weight 80, 84, 90

OML outer mold line 76, 101

RAE Royal Aircraft Establishment 48, 61, 67, 68, 76

RANS Reynolds-averaged Navier–Stokes 4–6, 14, 20, 21, 24, 29, 30, 33, 36–38, 40

SA Spalart–Allmaras 20

SMT Surrogate Modeling Toolbox 83

SQP sequential quadratic programming 12

TACS Toolkit for the Analysis of Composite Structures 15, 86

TSFC thrust-specific fuel consumption xii, 16

uCRM undeflected Common Research Model 48, 50, 53

VLM vortex-lattice method 20, 23, 24, 31

XDSM extended design structure matrix 10

Abstract

Numerical optimization has been applied to wing design problems for over 40 years. Over the decades, the scope and detail of optimization problems have advanced considerably. At the present time, the state-of-the-art in wing design optimization incorporates high-fidelity modeling of the steady-state aeroelastic response of the wing at both on-design and off-design operating conditions. Reynolds-averaged solutions of the Navier–Stokes equations coupled with linear finite element analysis offer the highest fidelity modeling currently tenable in an optimization context. However, the complexity of implementing and cost of executing high-fidelity aerostructural optimization have limited the extent of research on the topic. The goal of this dissertation is to examine the general application of these tools to wing design problems and highlight several factors pertaining to their usefulness and versatility.

Two types of wing design problems are considered in this dissertation: refining and exploratory. Refining problems are more common in practice, especially for high-fidelity optimization, because they start from a good design and make small changes to improve it. Exploratory problems are intended to have liberal parametrizations predisposed to have significant differences between the original and final designs. The investigation of exploratory problems yields novel findings regarding multimodality in the design space and robustness of the framework.

Multimodality in the design space can impact the usefulness and versatility of gradient-based optimization in wing design. Both aerodynamic and aerostructural wing design problems are shown to be amenable to gradient-based optimization despite the existence of multimodality in some cases. For example, a rectangular wing with constant cross-section is successfully converted, through gradient-based optimization, into a swept-back wing with transonic airfoils and a minimum-mass structure. These studies introduce new insights into the tradeoff between skin-friction and induced drag and its impact on multimodality and optimization. The results of these studies indicate that multimodality is dependent on model fidelity and geometric parametrization. It is shown that artificial multimodality can be eliminated by improving model fidelity and numerical accuracy of functions and

derivatives, whereas physically significant multimodality can be controlled with the application of geometric constraints.

The usefulness of numerical optimization in wing design hinges on the ability of the optimizer to competently balance fundamental tradeoffs. With comprehensive access to the relevant design parameters and physics models of the aerostructural system, an optimizer can converge to a better multidisciplinary design than is possible with a traditional, sequential design process. This dissertation features the high-fidelity aerostructural optimization of an Embraer regional jet, in which simultaneous optimization of airfoil shape, planform, and structural sizing variables yields a significantly improved wing over the baseline design. For a regional jet, it is shown that the inclusion of climb and descent segments in the fuel burn computation has a significant impact on the tradeoff between structural weight and aspect ratio. Another study addresses the tradeoff between cruise performance and low-speed, high-lift flight characteristics. A separation constraint at a low-speed, high-lift condition is introduced as an effective method of preserving low-speed performance while still achieving significant fuel burn reduction in cruise.

Chapter 1

Introduction

Wing design is a complicated endeavor. To be viable in a practical application a wing must generate lift, but do it stably and efficiently. Every contour of its surface plays a role, for better or for worse, in overall aerodynamic performance. It must withstand critical load conditions, yet any nonessential structural weight only detracts from its effectiveness. Additionally, a wing is generally expected to perform well at a variety of flow conditions in spite of aeroelastic deformations. All of these reasons and more give rise to a wonderfully intricate design problem, one that stymied dreamers for millennia and continues to demand significant effort in the modern era. At present, state-of-the-art wing design involves using numerical optimization in conjunction with high-fidelity computational models in order to determine the optimal wing for a set of requirements. These tools enable the human designer to spend less time manually iterating and instead focus on asking the right questions and interpreting results. However, high-fidelity wing design optimization is still relatively new and industry has a long history of success with more traditional design methods. The responsibility of demonstrating and promoting its value rests with academia and researchers. The purpose of this dissertation is to demonstrate the value of using gradient-based optimization to explore the design space in aerostructural wing design problems. Some of the questions examined include

- Is the design space multimodal? What are the implications of multimodality for a gradient-based optimizer?
- How does model fidelity affect the result of the optimization? What is the appropriate model for a given design problem?
- How do different design parameters influence wing performance? Does simultaneous optimization of all parameters produce a different design than a more traditional design method?

Given the open-ended nature of these questions, I do not expect to provide comprehensive answers that would be applicable to all cases. Rather, for the three case studies presented herein, the aim is to provide a template and encourage designers to use it to answer their own questions.

1.1 A Brief History of Wing Design

In 1799, Sir George Cayley first conceived of the fixed-wing airplane, brilliantly decoupling the lift problem from the thrust problem¹ in contrast to the bioinspired flapping flight concepts of his day. The following century saw numerous attempts at powered flight end in failure due to a collective preoccupation with muscling the aircraft into the air with more powerful (and heavier) engines. In contrast, the Wright brothers (among others) found success by taking a more holistic approach to aircraft design. They rigorously tested hundreds of airfoil shapes, leading to aerodynamically efficient wings and propeller blades. They also selected an efficient structural design which enabled a higher aspect ratio for their *Wright Flyer* than they had previously used on their gliders². Finally, they mastered the control of their aircraft through the novel use of wing warping. Individually, any one of these improvements may not have been sufficient, but taken altogether, they produced a step change in aircraft performance. Since that first success, the integrated nature of the aircraft has been one of its defining characteristics.

In the early days of aviation, experimental testing was the best way to build understanding of the aircraft and determine how to improve its performance. This was partly due to the fact that key aspects of aerodynamic theory were still undiscovered. Experimental testing was well-suited to the kinds of basic parameter sweeps being investigated at that time, however it is not an efficient tool for an iterative design process. It did not take long for aerodynamicists to uncover important new concepts—the Kutta-Joukowski theorem (1906), Prandtl’s boundary layer (1904)³ and lifting line (1918)⁴ theories, and Munk’s thin airfoil theory (1922)⁵ to name a few—that greatly enhanced design intuition. These theoretical breakthroughs were too general to be useful in detailed design comparisons, but they paved the way for more practical numerical methods that could have a place in the design process.

The gap between theory and practical design was bridged with the introduction of the National Advisory Committee for Aeronautics (NACA) family of airfoils by Jacobs in 1931⁶. The simplicity of the NACA airfoil definition, combined with the exhaustive empirical data published for a variety of camber and thickness distributions, made the NACA airfoils an invaluable tool in the wing design process. Around the same time, Theodorsen⁷ generalized the conformal transformation (originally developed by Joukowski) to allow the calculation of pressure distributions for air-

foils of any arbitrary shape. It would take another 30 years before methods were developed (and computers were capable enough) to calculate pressure distributions over arbitrary three-dimensional bodies. These so-called panel methods were first introduced for two-dimensional and axisymmetric shapes by Smith and Pierce⁸. Hess and Smith⁹ published the generalized method for three-dimensional bodies in 1967. Although the accuracy of inviscid panel methods would be surpassed in short order, they continue to be used in aircraft design to the present day. In later years, Hess explained the reason for their enduring usefulness: “Even when their results fail to give the proper experimental values, they are frequently useful in predicting the incremental effect of a proposed design change or in ordering various designs in terms of effectiveness”¹⁰. This statement underscores the value of computational fluid dynamics (CFD) in the design process, regardless of its fidelity to reality. This is not to say that a very low-fidelity model should be used in detailed wing design. Rather, for every legitimate aerodynamics model, there is a subset of the design space over which the model will provide useful information from which design decisions can be made. Critically, the designer must still determine at what stage of the design process and with respect to which design parameters a given model fidelity is useful.

Regardless of the method used to analyze the design, there remains the fundamental issue of selecting the design parameters for a given iteration of the design cycle. Ideally, the designer should have some intuition regarding the tradeoffs in the design space and the sensitivity with respect to different parameters. This is a tall order even for experienced designers, yet the alternative is basically educated guessing. Richard Whitcomb famously spent hours in the wind tunnel with a putty knife, painstakingly modifying the wing section until arriving at the supercritical airfoil. One solution is to invert the problem and find the geometry that corresponds to a prescribed pressure distribution. This inverse design method was pioneered by Jacobs who, despite his great success with NACA airfoils, came to the realization that “continued empirical testing of airfoils, no matter how systematic, would not lead to greatly improved shapes except by luck”⁶. Mangler¹¹ and Lighthill¹² independently worked on solving the same problem, and modern inverse airfoil design methods derive from their efforts. However, inverse design methods presuppose a knowledge of the optimum pressure distribution, which is not generally the case in design problems. The solution to these shortcomings in the design process came in the form of numerical optimization.

1.2 Numerical Optimization of Wings

The solution of general engineering design problems via numerical optimization originated in structural mechanics. Prior to 1960, optimal structural design was carried out with specialized formulations that would not have been generalizable to a broader class of problems. Schmit¹³ recognized that engineering design problems were naturally suited to be framed as general nonlinear optimization problems. Schmit's innovation was motivated by the desire to solve more general structural design problems (e.g. multiple load conditions, custom objective functions, etc.), but in generalizing the methodology, he also opened the door to numerous other engineering applications. Schmit called his method "structural synthesis", which aptly conveys the fusion of various research fields which is characteristic of design optimization.

The first application to airfoil design was the result of a collaboration between a former student of Schmit and two National Aeronautics and Space Administration (NASA) aerodynamicists. Hicks et al.¹⁴ demonstrated gradient-based optimization on airfoils in inviscid flow using the transonic small-disturbance equations. They were motivated by the need for a better automated design method—one that was simple to implement, generalizable to different flow regimes and three dimensions, and capable of handling generic constraints. In 1978, Hicks and Henne¹⁵ extended the method to optimize the design of a three-dimensional wing using a full potential inviscid aerodynamics code. Regarding their work, they stated: "The primary objective of this investigation was to demonstrate the usefulness and versatility of numerical optimization in wing design." In three optimization problems, they used the eponymous Hicks–Henne bump functions to optimize a wing for various objectives and constraints. Thus, more than 40 years ago, the primitive form of wing design optimization emerged. The intervening years have seen remarkable advances in computing capability, modeling fidelity, and optimization efficiency, all of which have improved the usefulness and versatility of numerical wing design. Yet the basic idea remains the same as that presented by Hicks and Henne.

Numerical design optimization has benefited immensely from progress in general computing and numerical analysis. The fidelity of CFD tools increased as panel methods gave way to solutions of the full potential equations, the Euler equations, and finally the Reynolds-averaged Navier–Stokes (RANS) equations. Each of these advances enabled more accurate predictions of the complex physical phenomena inherent to transonic flow conditions. The scope and complexity of problems that can be realistically analyzed has also increased with the exponential rise of computer processing power and memory limits. At present, second-order finite volume RANS solvers are widely used in industry to complement and gradually replace wind tun-

nel testing. Modeling accuracy is of paramount importance to design optimization because it limits the degree of detail that can be admitted in the parametrization. A good aerodynamic design is often determined by small or nuanced design features, yet it is pointless to give the optimizer control of these features if they are not accurately represented in the physical model.

At the same time, the cost of the optimization problem mounts with higher model fidelity and greater numbers of design variables. This rising cost can be alleviated by using gradient-based optimizers, which generally require fewer function evaluations to converge to an optimum, thus reducing the total computational time. However, the cost of naïvely computing derivatives by finite differencing also scales with the number of design variables. The fundamental innovation that paved the way for high-fidelity aerodynamic shape optimization was the adoption of the adjoint method from control theory by Pironneau. He applied adjoint theory to obtain minimum drag profiles for Stokes flow¹⁶ and the incompressible Euler equations¹⁷. Later, Jameson¹⁸ derived the adjoint of the compressible Euler equations and used it to optimize airfoils¹⁹ and wings²⁰. The cost of using the adjoint approach to compute derivatives is independent of the number of design variables. Hicks and Henne¹⁵ used finite differences to calculate derivatives and were limited to 11 design variables in their optimizations. By contrast, Jameson²⁰ optimized the positions of up to 4,224 surface mesh points in one of his first forays into wing optimization. While this was an impressive demonstration of the scalability of the adjoint method, it is far from being the ideal way to parametrize a wing.

Computer-aided design (CAD) software is a natural first choice for design parametrization because it allows for intuitive design variables and is ubiquitous in industry. In fact, many modern CAD packages can be bundled into comprehensive computer-aided engineering (CAE) frameworks that often include optimization capabilities. However, CAD has not yet been widely adopted for aerospace optimization applications because it has not been adapted to work with the adjoint method. Many researchers have approximated CAD derivatives using finite differences, but this approach is untenable for more than a handful of design variables due to computational cost. Instead, more nimble methods have been adopted for geometric parametrization in wing design. These methods provide more intuitive control of the geometry than merely displacing mesh points, yet offer efficient gradient calculation and negligible computational overhead. Samareh²¹ gave a comprehensive review of such methods.

Aerodynamic shape optimization (ASO) lies at the intersection of numerical optimization, CFD, and CAD. The fruits of ASO are too numerous to be comprehensively reviewed here, but I will highlight some of the key milestones. Reuther et al.²² extended the applicability of adjoint optimization to large-scale problems

with full-configuration, multiblock CFD meshes through improvements to the design parametrization method, mesh perturbation scheme, and parallel implementation^{23,24}. The development of the RANS adjoint^{25–27} enabled the realization of more practical designs, especially in the transonic regime. It was later shown that significant differences arise between Euler and RANS-optimized wings^{28,29}. ASO has been applied to conventional^{30–32} and unconventional aircraft^{33–35}. It has proven useful for exploring nonplanar wings^{36,37} and wingtip devices^{38,39}. The development of a set of benchmark optimization problems by the AIAA Aerodynamic Design Optimization Discussion Group (ADODG) was an important step in establishing expectations and standards for ASO. For a more thorough review of the history of ASO, see Jameson and Ou⁴⁰, and for a deeper look at recent advances, see Skinner and Zare-Behtash⁴¹.

Wing design is inherently a multidisciplinary problem, with the two dominant disciplines being aerodynamics and structures. Any design that is produced by ASO either blatantly violates structural requirements or is handicapped by overly conservative constraints which are put in place to prevent such violations. An optimizer needs access to rich information from both aerodynamic and structural models in order to optimally balance fundamental tradeoffs in the wing design problem. The structural optimization community was first to consider multidisciplinary effects in wing design. Haftka⁴² optimized a wing structure subject to flutter constraints modeled with second-order piston theory aerodynamics. Soon thereafter, Haftka⁴³ combined a lifting line model with a simple finite element model to perform one of the earliest aerostructural optimizations. But at these early stages, the prospect of analyzing—let alone optimizing—a high-fidelity model considering both aerodynamics and structures was out of the question. An eminent aircraft designer of that era, Dietrich Küchemann, presciently wrote, “This should be one of the aims for the future: we want an integrated aerodynamic and structural analysis of the dynamics of the flying vehicle as one deformable body, and to use that for design purposes”⁴⁴. This yearning for more integrated analysis of aircraft wings was actually part of the more expansive multidisciplinary analysis and optimization (MDAO) movement, which is covered in more detail by Martins and Lambe⁴⁵.

Küchemann’s dream of a high-fidelity aerostructural design optimization framework began to come to life at the turn of the 21st century. With increasingly sophisticated CFD and computational structural mechanics (CSM) tools came the need for transfer schemes^{46–48} that could accurately represent the fluid-structure interaction (FSI). The development of efficient aerostructural solvers and methods for multidisciplinary derivative computation opened the door to high-fidelity multidisciplinary design optimization (MDO). Maute et al.⁴⁹ optimized a simple swept wing using Euler CFD and structural finite element analysis (FEA), but they only used the direct

method to calculate derivatives and were thus limited in the number of design variables. Martins et al.⁵⁰ developed the coupled adjoint for aerostructural systems and applied it to the optimization of a supersonic business jet⁵¹. Since then, many researchers have employed aerostructural optimization using the RANS equations and complete finite-element wingbox descriptions^{52–56}. This advancement in the modeling fidelity used in optimization is only possible with efficient parallel implementation, adjoint sensitivity calculation, and rapid, robust solvers. However, even when all of these details are considered, a high-fidelity optimization can be quite computationally expensive.

1.3 Motivation and Objectives

The preceding two sections have narrated the evolution of wing design, culminating in the development of MDO. Currently, the aircraft industry is in the process of incorporating MDO into its standard design cycle. However, the transformation is by no means complete or even fully supported. Borrowing a phrase from Hicks and Henne¹⁵, the objective of this dissertation is to further validate the “usefulness and versatility of numerical optimization in wing design.” The words *useful* and *versatile* are unexpectedly fitting for the content covered herein. The following paragraphs expand on how I demonstrate usefulness and versatility in this dissertation.

Versatility implies robustness, generality, and flexibility. Optimization is touted as an invaluable tool for studying and designing unconventional aircraft for which there is no historical data to inform design decisions. A versatile MDO framework should adapt easily to any aircraft configuration and explore the design space to find the optimal design. Gradient-based optimizers are often criticized in this respect because, by construction, they converge to a single local minimum, which may or may not be the global optimum. This criticism carries the implication that MDO is better suited for refining already good designs rather than exploring novel design spaces. Many researchers have studied exploratory optimization problems to determine whether this is a valid claim. Airfoil optimization problems have been shown to be unimodal and well-disposed to solution by a gradient-based optimizer⁵⁷. Likewise, optimization of the local airfoil shape of three-dimensional wings (fixed planform) is a unimodal problem^{58–60}

When planform variables are added to the design problem, multiple local minima do appear in the design space^{57,61,62}. However, it would be irresponsible to conclude from these results that MDO is not appropriate for these types of problems. A more conscientious study of the fundamental reasons for multimodality is warranted. In some cases, the local minima are so close to each other that the differences are insignificant, given the modeling tolerances⁵⁸. Alternatively, the

designs might be quite different but the objective can be very similar due to a very flat design space. Sometimes local minima may exist due to known physical phenomena (e.g. forward and aft swept wings). Local minima might also exist when the modeling fidelity is inadequate to distinguish between two designs that have very slight differences. In this dissertation, I study two different exploratory optimization problems with the goal of elucidating the fundamental reasons for the existence of multiple local minima.

A useful optimization framework is one that can readily be applied to practical design problems (as opposed to academic studies). The usefulness of MDO in wing design has already been well demonstrated in the literature. In this dissertation, I address a few of the practical aspects of wing design optimization. With regards to the multimodality issues previously discussed, I propose specific design constraints to eliminate or mitigate multiple local minima from the design space. Throughout the dissertation I also provide practical tips for optimization practitioners, such as the use of successive optimization runs with increasingly finer grids to reduce overall computational cost.

One of the common pitfalls of wing optimization is the tendency for the optimizer to maximize performance at specific design conditions at the expense of robust performance across the entire flight envelope⁶³. This problem is frequently addressed by including the performance at multiple design conditions in a single, composite objective function. In the literature, however, multipoint optimization is usually aimed at improving robustness in cruise flight, where the aircraft spends most of its time^{24,58,64}. In this dissertation, I show that a wing optimized for multiple cruise flight conditions still may not be flightworthy at low-speed, high-lift conditions. I experiment with using a separation constraint to ensure good low-speed performance on a cruise-optimized wing design.

The capstone optimization problem of this dissertation involves the redesign of an Embraer regional jet wing. For this short-haul aircraft, I show the importance of considering the performance of the entire flight profile, including climb, cruise, and descent segments, rather than just optimizing for cruise performance as is often done in the literature for larger aircraft. The final results of this optimization problem are a definitive demonstration of the benefits of simultaneously optimizing twist distribution, cross-sectional shape, planform, and structural sizing in order to arrive at the best wing design.

In all of these studies, our approach is to use optimization to extract design insights and understand tradeoffs. Philosophically speaking, I do not see MDO as a threat to the human designer but rather as a tool to improve the efficiency and innovative potential of the design process.

1.4 Contributions and Outline

The original contributions in this dissertation are listed here:

1. Developed new insights into the tradeoff between viscous and inviscid drag in relation to optimization and multimodality.
2. Demonstrated practical ways to reduce multimodality in the design space.
3. Demonstrated plank to transonic wing with aerostructural optimization.
4. Pioneered low-speed separation constraint to improve off-design robustness.
5. Investigated impact of mission analysis method on optimization of short-haul aircraft.
6. Demonstrated industrial application of MDO on Embraer regional jet.

The dissertation is organized as follows: Chapter 2 introduces the MDO of aircraft configurations with high fidelity (MACH) framework, which has been developed at the University of Michigan for the purpose of high-fidelity aircraft design optimization. Since the tools in the MACH framework are used throughout the dissertation, I opt to introduce them all at once rather than repeatedly in each succeeding chapter. In Chapter 3, I study an exploratory ASO benchmark problem and the multimodality present in its design space. In Chapter 4, I extend our study of exploratory wing optimization to the aerostructural design space. I also experiment with various methods for producing robust optimal designs, especially with regards to off-design, low-speed performance. Chapter 5 deals with the optimization of the Embraer regional jet wing and the study of the impact of mission analysis on the optimized wing design. Finally, in Chapter 6 I draw conclusions, restate contributions, and propose future avenues of research.

Chapter 2

Overview of the MACH Framework

The MACH framework offers an automated approach to aircraft design⁶⁵. At its core is a coupled aerostructural analysis and adjoint capability, but MACH also encompasses numerous other modules that contribute to the overall goal of automated aircraft design. The components of MACH used in this dissertation are depicted in the extended design structure matrix (XDSM)⁶⁶ diagram in Figure 2.1. The XDSM diagram highlights both the component hierarchy and the data (grey lines) and process (black arrows) flows. In particular, this XDSM diagram shows how the framework handles the design optimization problem in Chapter 5, which incorporates aerostructural and mission analyses. All of the MACH modules are implemented in the Python programming language, although many of them wrap compiled code for more computationally heavy tasks. The framework is designed to be lightweight and modular, so that simpler problems can be solved by only including the necessary modules. For instance, an aerodynamic analysis can be carried out with a simple Python script that only imports a subset of the modules shown in Figure 2.1. This chapter provides a brief introduction to each of the modules used in the work presented in this dissertation.

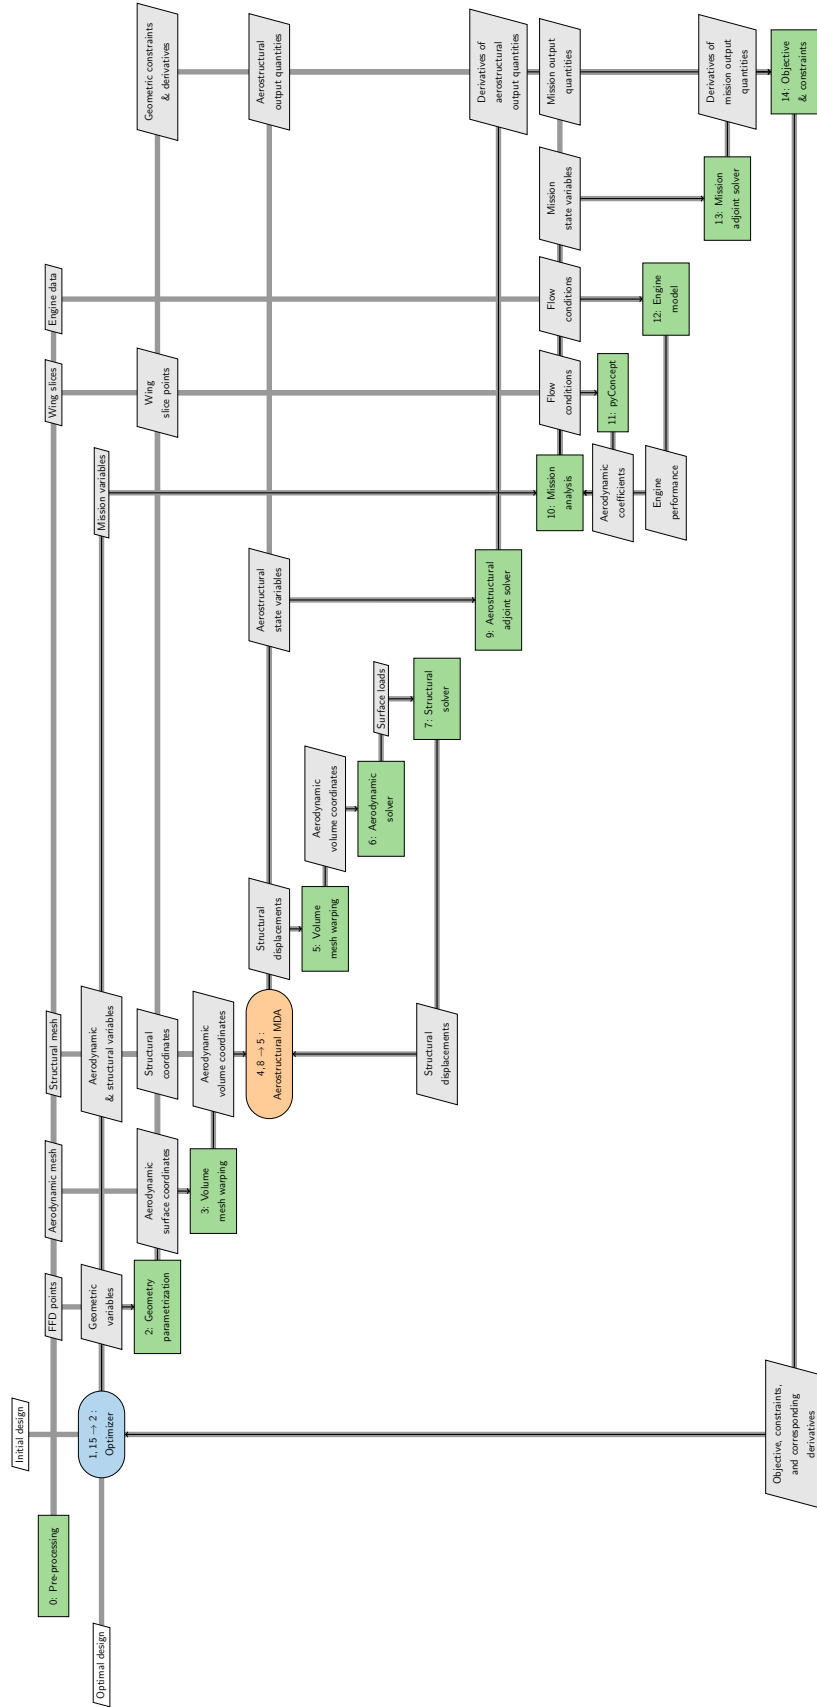


Figure 2.1: XDSM diagram of aerostructural optimization with MACH.

2.1 Optimization Algorithm

The framework is optimizer-agnostic and has been tested with a variety of different optimizers⁶⁰. We typically use pyOptSparse* because it provides a common interface to a variety of different optimizers, allowing the user to switch the optimizer without having to reprogram the problem definition. In this dissertation we use SNOPT⁶⁷ (wrapped with pyOptSparse) exclusively. SNOPT is a sequential quadratic programming (SQP) optimizer which has been used extensively for both aerodynamic shape optimization and aerostructural design optimization.

2.2 Geometry

2.2.1 Geometric Parametrization

The geometry is parametrized with a free-form deformation (FFD) scheme⁶⁸ implemented in pyGeo[†] by Kenway et al.⁶⁹. The FFD formulation is attractive because it parametrizes the geometric changes rather than the geometry itself, enabling control of objects for which the underlying geometric definition is unknown. In pyGeo, an FFD volume is a structured three-dimensional grid of B-spline control points. The process of embedding a set of points in the FFD volume involves executing a Newton search to find the parametric position of each point in the B-spline reference space. Any number of point sets can be embedded in the FFD volume. As the positions of the FFD control points change, the embedded point sets deform continuously according to the B-spline mapping. Conveniently, the derivative of an embedded point with respect to the B-spline control points is defined (and computed) analytically.

All that remains is to decide how to control and coordinate the displacement of the control points. In the pyGeo implementation, the design variables can be set up to enact local or global changes of the embedded geometry. The global design variables allow the user to coordinate the movement of multiple FFD control points to create large-scale deformations. This coordination is facilitated by a user-defined B-spline reference axis curve. Transformation properties such as translation, rotation, and scaling design variables can be specified at each reference axis control point. The B-spline representation of the reference axis ensures a smooth variation of these transformation properties between the control points. Each FFD control point is linked to a parametric position along the reference axis and assumes the transformation properties corresponding to that point of attachment. The various points used in the parametrization scheme are indicated in Figure 2.2. The user can

*<https://github.com/mdolab/pyoptsparse>

†<https://github.com/mdolab/pygeo>

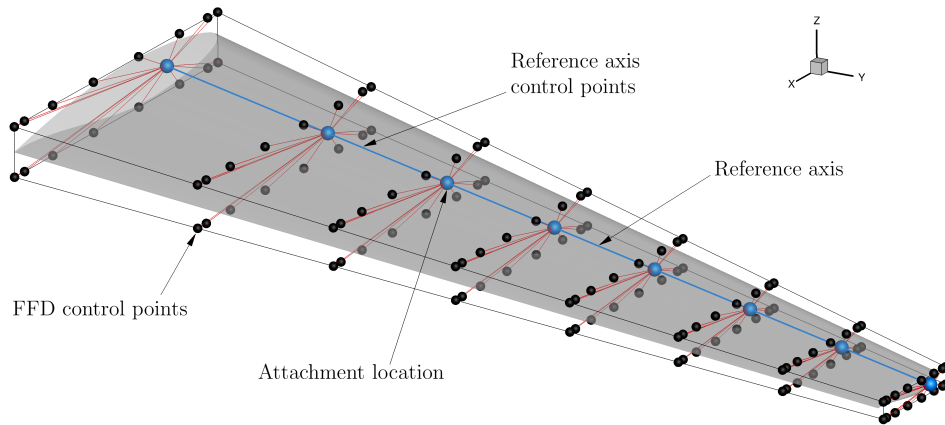


Figure 2.2: The geometric parametrization uses FFD control points for local control and a B-spline reference axis for global control.

create custom functions to specify how the transformation properties are applied to the reference axis control points. When multiple design variables affect a given control point, the operations are combined linearly. In general, a global design variable will produce a nonzero derivative for multiple control points. The derivatives of the FFD control points with respect to the global design variables are computed using the complex-step method.

The local design variables control the displacement of individual control points for precise shape modifications. The user can choose whether the displacement is along a global coordinate axis (x , y , or z) or along a local reference axis defined by the FFD grid. In the latter case, the displacement direction is dependent on any rotations imposed by the global design variables.

2.2.2 Geometric Constraints

In addition to directly choosing the type, number, and bounds of geometric design variables, the user can also apply geometric constraints to restrict the design space. These constraints are often used to prevent design changes that would violate requirements that are not considered in the optimization problem. For instance, thickness constraints are used in ASO to prevent thickness decrease that would otherwise be prevented by structural constraint violations if the structure was modeled. The geometric constraints in the MACH framework are points-based. Each constraint is defined by one or more points which are embedded in the FFD volume. For example, a thickness constraint is defined with two points and volume constraint is defined with a grid of points making up a number of hexahedra. Some of the available constraints are depicted in Figure 2.3. Updates to the point coordinates and their derivatives are calculated inexpensively with pyGeo. The values and derivatives of the constraints are easily calculated with respect to their

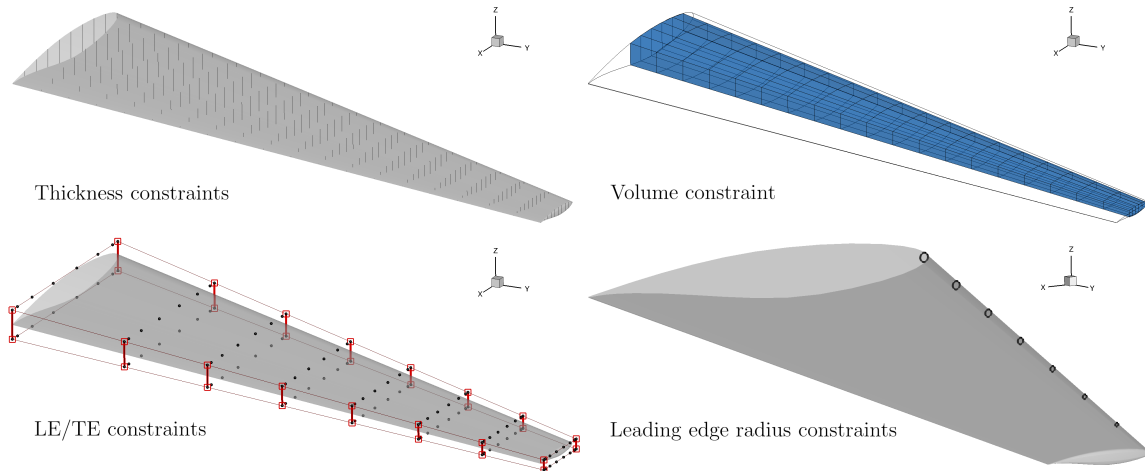


Figure 2.3: Various point-based geometric constraints are possible in the MACH framework. Linear constraints can also be added to the FFD control points.

constitutive point sets.

In addition to the point-based constraints, pyGeo also offers functionality to create linear constraints on the FFD control points. For example, the LE/TE constraint forces the leading edge and trailing edge control points to move in equal and opposite directions to prevent shearing twist. Without this linear constraint, there would be a dependency between the twist variables and the shape variables, which is generally undesirable in an optimization context.

2.3 High-fidelity Physics Modeling

2.3.1 Aerodynamic Analysis

The flow solver in MACH is ADflow[‡]. ADflow is a finite-volume CFD solver for cell-centered multiblock and overset meshes. It solves the compressible Euler, laminar Navier–Stokes, and RANS equations with a second-order accurate spatial discretization. The solver employs a variety of numerical methods to converge to a steady-state solution, including multigrid, approximate Newton–Krylov (ANK), and Newton–Krylov (NK) algorithms. The combination of these various iterative methods makes ADflow robust and fast⁷⁰. ADflow also solves the discrete adjoint equations, enabling efficient computation of derivatives independent of the number of design variables. The solution of the discrete adjoint in ADflow relies on the ADjoint^{71,72} approach, which uses automatic differentiation (AD) to compute partial the derivatives and a Krylov method to solve the linear system.

[‡]<https://github.com/mdolab/adflow>




Figure 2.4: Representative depiction of smeared stiffness model.

2.3.2 Mesh Warping

The mesh used for solving the fluid dynamics equations must be representative of the geometry as it changes throughout the optimization process. Rather than regenerating the mesh at each iteration, we modify the coordinates of the original mesh to make them consistent with the updated geometry in a process known as *mesh warping*. In the setup phase, the surface coordinates of the mesh are embedded in the FFD volume. At each iteration, the updated surface coordinates are sent to the mesh warping algorithm, IDWarp[§]. IDWarp uses an inverse-distance weighting method to update the volume mesh coordinates based on the changes to the surface mesh⁷³. The derivatives of the volume mesh nodes with respect to the surface mesh nodes are computed using automatic differentiation (AD).

2.3.3 Structural Analysis

MACH uses the Toolkit for the Analysis of Composite Structures (TACS)[¶] to model the structural mechanics of wings. TACS was developed by Kennedy and Martins⁷⁴ specifically for the analysis and optimization of the thin-walled structures typical in aircraft. Typically, we analyze wingbox structures comprising ribs, spars, and skins. Each part of the structure is discretized with quadrilateral shell elements, which are grouped into rectangular panels. In the simplest case, the panels are modeled with isotropic constitutive properties. Alternatively, virtual stiffeners can be incorporated into the panel stiffness matrix to achieve the effect of a blade-stiffened panel (see Figure 2.4) Each panel can have its own variables to control panel thickness, stiffener thickness, stiffener height, and stiffener pitch.

The structural model is used to predict stress and buckling limits which can be used as constraints in an optimization problem. Rather than constrain the structural failure on an element-by-element basis, we aggregate the failure constraints using the Kreisselmeier–Steinhauser (KS) function. The KS is a conservative approximation of the maximum, meaning that it will overestimate the failure of the structure and provide us with a conservative design. Without constraint aggregation, there would

[§]<https://github.com/mdolab/idwarp>

[¶]<https://github.com/gjkennedy/tacs>

be thousands of structural constraints, each of which would require an adjoint calculation. To make the adjoint approach worthwhile, it is necessary to have far fewer function outputs than design variables.

2.3.4 Aerostructural Coupling

The structural node displacements and aerodynamic surface loads are transferred between the aerodynamic and structural meshes using a rigid link, load and displacement transfer scheme first introduced by Brown⁴⁷ and subsequently described in the context of MACH^{65,75,76}. We obtain the solution of the aerostructural system with a Gauss–Seidel iterative scheme. A Krylov method is used to solve the coupled adjoint of the multidisciplinary system.

2.4 Mission Analysis

Often the objective of an aircraft design optimization problem is to decrease direct operating cost (DOC). The cost of fuel makes up a significant portion of DOC. Additionally, aircraft original equipment manufacturers (OEMs) and airlines are increasingly pressured to reduce the detrimental effects of fuel burn on the environment. Thus, excess fuel burn comes at a premium both financially and environmentally.

There are various methods to estimate fuel burn computationally. For preliminary takeoff weight estimation, the fuel burn for a given segment of the mission is reasonably predicted using fuel fractions extracted from historical data. However, for the purposes of optimization it is necessary to derive a formula relating the fuel burn to the actual design of the aircraft. The Bréguet range equation is a simple yet mathematically sound relationship between the aerodynamic, structural, and propulsion performance of the aircraft. For a jet aircraft the rate of weight change is proportional to thrust

$$\frac{dW}{dt} = -T c_T \quad (2.1)$$

where c_T , the thrust-specific fuel consumption (TSFC), is dependent on the engine model, Mach number, and altitude. When divided by the instantaneous flight speed, this can be transformed into an integral equation to compute range.

$$\begin{aligned} \frac{\frac{dW}{dt}}{\frac{ds}{dt}} &= \frac{-T c_T}{V} \\ \frac{ds}{dt} &= \frac{V}{-T c_T} \frac{dW}{dt} \\ R &= \int \frac{ds}{dt} dt = \int \frac{V}{-T c_T} dW \end{aligned} \quad (2.2)$$

For range covered in cruise, we can assume steady flight, meaning $T = D$ and $L = W$. Additional assumptions of constant flight speed, TSFC, and L/D yield the Bréguet range equation.

$$R = \frac{VL}{c_T D} \int \frac{-1}{W} dW$$

$$R = \frac{V}{c_T} \frac{L}{D} \ln \frac{W_2}{W_3} \quad (2.3)$$

In this equation, W_2 and W_3 refer to the aircraft weight at the beginning and end of cruise, respectively. Although Equation 2.3 bears the name of French aviation pioneer Louis Bréguet, according to Cavcar⁷⁷ it was independently derived by Devillers⁷⁸ in 1918 and Coffin⁷⁹ in 1920. Equation 2.3 can be further rearranged to calculate fuel burn as follows:

$$\frac{W_2}{W_3} = \exp\left(\frac{c_T R D}{V L}\right) \quad (2.4)$$

This equation succinctly relates the three dominant disciplines affecting fuel burn, enabling the optimizer to balance them appropriately. However, due to the stated assumptions, Equations 2.3 and 2.4 apply specifically to a cruise-climb flight mode, in which altitude steadily increases to enable the constant L/D assumption.

For a more general approach to range and fuel burn estimation, MACH framework contains a module called MissionAnalysis. Liem et al.⁸⁰ created MissionAnalysis to enable surrogate-based mission analysis within the context of an aerostructural optimization problem similar to the one we are considering here. More details on the approach used are given in their paper. Although the basic function of MissionAnalysis remains the same, we made some significant changes to improve its generality and practicality. Originally, the code was written with internal surrogate models for the aerodynamics (C_L , C_D , and C_{M_y}). This approach required the creation of aerodynamic data sets with respect to altitude, Mach number, angle of attack, and tail angle at every optimization iteration. To make the code more general, we replaced the internal surrogate model with a set of callback functions through which the user can provide aerodynamic and engine data to MissionAnalysis. Derivatives of the functions of interest with respect to external design variables are also supported in the code. Derivatives in MissionAnalysis are computed using automatic differentiation to achieve machine precision.

In MissionAnalysis, the various mission segments are broken up into integration intervals. The fuel burn, distance, and elapsed time for each interval are computed using integral equations similar to Equation 2.2. The states of the mission model system are the weights at the interval endpoints. A solution to the nonlinear system is obtained by driving the weight differences between endpoints of consecutive intervals to zero. As in Liem et al.⁸⁰, a line-search stabilized Newton's method

is used to solve the nonlinear system and the Jacobian is calculated using finite differences. For our purposes in this dissertation, the solver is set up to perform in the following manner. The user or optimizer provides the solver with a fixed weight, an initial fuel weight, and a cruise altitude. The solver converges the residuals to zero to reach a valid state for the prescribed inputs. The outputs of the solver are the range, duration, and fuel burn over the course of the mission. Incidentally, the input fuel weight and the output fuel burn can be different if the aircraft starts with more than enough fuel to complete the given mission. With this setup, it is possible for the optimizer to choose inputs that lead to an impractical solution. For example, if the drag at a given altitude is greater than the maximum available thrust, the solution will be invalid. Therefore, in addition to the standard outputs, MissionAnalysis computes slack functions that can be used to force a valid solution. For the climb segments, the slack function at each node along the mission profile is the actual climb rate minus the required residual climb rate. For the cruise and descent segments, the slack function is the available thrust minus the required thrust. In an optimization problem, these slack functions are constrained to be greater than zero.

2.4.1 Conceptual-level Aircraft Model

pyConcept is a module in the MACH framework that can be used to calculate geometric parameters such as reference area, mean aerodynamic chord, and sweep angle. The wing geometry is represented by points located along the leading and trailing edges and the location of maximum thickness on the baseline geometry. These points are extracted from airfoil slices of the baseline wing and then embedded in the FFD. Geometric parameters are then dependent on the FFD design variables via the embedded points.

pyConcept also provides estimates of aerodynamic performance based on fundamental geometric properties. For instance, induced drag can be approximated as

$$C_{D,i} = \frac{C_L^2}{\pi AR e} \quad (2.5)$$

where both AR and e can be calculated solely from geometric properties and flow conditions. Similarly, parasitic drag is computed using the component buildup method suggested by Raymer⁸¹, wherein the drag of each component is a function of wetted area, the flat-plate skin friction coefficient, a form factor, and an interference factor. These estimates are not expected to yield accurate predictions, but are intended to capture the trends due to changing geometric parameters. The novelty of this approach is that we are able to maintain the FFD parametrization as a means of adjusting the geometric parameters relevant to the conceptual formulas. Like

everything else in the MACH framework, the fundamental philosophy in developing this capability was to allow for efficient, accurate derivative computation. The derivatives of the pyConcept functions with respect to the geometric parameters are analytically derived and the derivatives of the geometric parameters with respect to the airfoil points are computed with the complex-step method. The derivatives of the points with respect to the design variables are already handled in pyGeo, and in the end, derivatives of the pyConcept functions with respect to the design variables are obtained by combining the three different sets of gradients using the chain rule.

Chapter 3

Exploration of Aerodynamic Wing Design

This chapter reports on a study of the sixth benchmark optimization problem from the AIAA Aerodynamic Design Optimization Discussion Group (ADODG)*. The purpose of this benchmark problem is to study multimodality in the design space of a subsonic, rectangular wing. This test case differs from the Common Research Model (CRM) wing optimization problem in that it gives the optimizer much more freedom to modify the shape of the baseline wing. Practical design constraints are neglected in favor of liberating the optimizer to explore the entire design space and perhaps discover a novel optimal design for the prescribed conditions. We use this test case as a foundation for a series of optimization studies in which we explore the physical reasons for multimodality in aerodynamic wing design. In Section 3.1, we introduce the test case and the tools used to perform the optimizations. Before reporting the results of the full test case, we investigate in Section 3.2 the implications of including each design variable with some preliminary optimization studies. Then we analyze the full case and provide some concluding commentaries in Sections 3.3 and 3.5, respectively. In this work, our purpose is not necessarily to do an exhaustive search of the design space to find all possible local minima, or even to rate this problem in terms of its multimodality. Other researchers have done excellent work in addressing these goals^{61,82}. Our aims are to show whether or not multiple local minima exist, clarify the factors influencing multimodality in the problem, and elucidate whether or not multimodality should be a concern for practical applications of MDO in aircraft design.

*<https://sites.google.com/view/mcgill-computational-aerogroup/adodg>

Table 3.1: Baseline geometry performance at $C_L = 0.2625$

Grid	Cells	α	C_D (cts)	$C_{D,v}$ (cts)	$C_{D,p}$ (cts)
Euler L3	180,992	3.023	42.694		
Euler L2	1,447,936	3.030	38.997		
RANS L2	306,432	3.206	144.013	76.378	67.636
RANS L1	2,451,456	3.205	130.545	69.702	60.843

3.1 Methodology

3.1.1 Multi-Fidelity Approach

The main goal of the ADODG Case 6 optimization problem is to study the existence of multiple local minima in the design space. In addition to this primary goal, we seek to understand whether such local minima reflect the real physics involved, or are merely artifacts of the modeling and discretization errors. In addressing these two goals, we found it useful to combine results from multiple sources of information. All told, we use three different physics models to analyze the aerodynamic performance of the wing: the RANS equations with a Spalart–Allmaras (SA) turbulence model, the compressible Euler equations, and a vortex-lattice method (VLM). The RANS and Euler equations are solved using ADflow^{28,83} and the VLM is implemented in OpenAeroStruct⁸⁴. The details of these solvers and their respective workflows are described in the following two sections.

3.1.2 Geometry description

The baseline wing geometry for the Euler and RANS analyses is planar with a chord of 1.0 m and a NACA-0012 airfoil cross-section. The wingtip cap is a perfect revolution about the airfoil chord line and adds 0.06 m to the 3.0 m rectangular portion of the wing, bringing the total semispan to 3.06 m. The Euler geometry has a sharp trailing edge, while the RANS geometry has a blunt trailing edge with a thickness of 2.52 mm. We generate the surface meshes using Ansys ICEM CFD and extrude the volume meshes using hyperbolical marching. The meshes are oriented with the x-axis in the streamwise direction, the z-axis out the wing, and the y-axis in the vertical direction. The quality of these meshes is tested in a grid convergence study at the nominal baseline condition ($M = 0.5$, $Re = 5 \times 10^6$, $C_L = 0.2625$), the results of which are plotted in Figure 3.1. The drag values converge nearly linearly as the number of cells to the power of 2/3 is increased, indicating asymptotic convergence. Table 3.1 lists the data for the baseline grids that are used in the optimization studies. For these meshes, x is the streamwise direction and y is the lift direction.




Figure 3.1: Grid convergence study at $M = 0.5$, $Re = 5 \times 10^6$ and $C_L = 0.2625$.

3.1.3 Wing parametrization

Our parametrization uses both global and local design variables, which are explained in detail in Section 2.2.1. Figure 3.2 depicts the FFD volume with the control points as black dots and the design variables definitions shown with color-coordinated arrows. The nominal FFD volume has nine spanwise control sections and 12 chordwise control points per section with half of the points on the upper surface and the other half on the lower surface. The global variables are linked to the displacement and rotation of axial control points (shown as red squares) along the reference axis. Each of these axial control points dictates the global movement of an entire FFD control section. For example, the blended winglet shown in Figure 3.2 was created by setting the dihedral variable with the following values:

$$x_{\text{dihedral}} = [0, 0, 0, 0, 0.03, 0.09, 0.24, 0.51, 1]$$

The user defined function for the dihedral variable displaces each of the nine axial control points in the y -direction according to the corresponding value in x_{dihedral} . Additionally, the dihedral function is set up to rotate the FFD sections so that they remain perpendicular to the reference axis. The local shape variables are handled differently. Each spanwise control section is assigned a unique reference frame with the section plane normal as the \hat{e}_k axis, the \hat{e}_i axis aligned with the streamwise direction, and $\hat{e}_j = \hat{e}_k \times \hat{e}_i$. The local shape variables control the movement of the control points along their respective \hat{e}_j axes. When the reference axis is displaced vertically, the control sections and their respective reference frames are automatically rotated to remain perpendicular to it. This behavior is depicted in the formation of the blended winglet on the wing in Figure 3.2. The displacement vectors of the local shape variables in the wingtip control section are rotated as the winglet forms to allow sectional control of the airfoil section of the winglet. This functionality en-

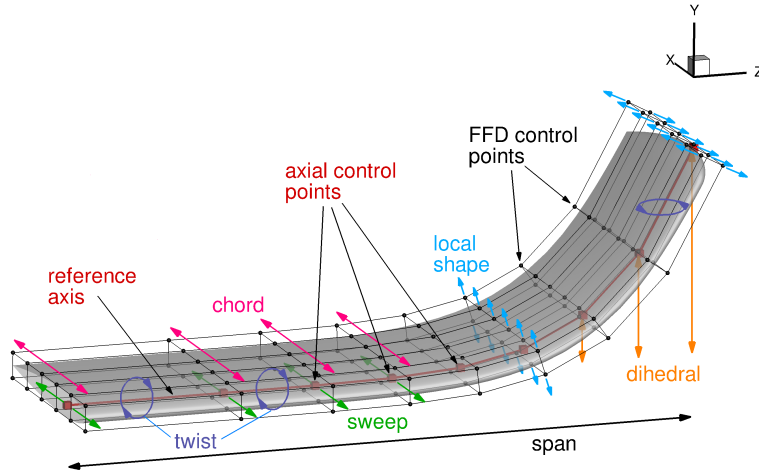


Figure 3.2: Geometric parametrization of the aerodynamic multimodality optimization problem. The wing surface is embedded in the FFD volume. Design variable definitions are indicated with arrows.

ensures that the wing surface does not shear, causing negative volumes, when large changes in dihedral are introduced.

3.1.4 Low-fidelity Optimization: OpenAeroStruct

OpenAeroStruct⁸⁴ is an open-source low-fidelity aerostructural optimization suite developed using the OpenMDAO framework⁸⁵. The aerodynamic analysis in OpenAeroStruct is performed using a VLM to compute induced drag and a modified flat-plate skin-friction drag approximation to estimate parasite drag. These low-fidelity models provide reasonable estimates at a low computational cost. A single analysis takes less than one second and a full optimization takes on the order of 10 seconds on a single processor. We use this low-fidelity code because it allows us to explore the design space quickly and gain insights that can be further explored with high-fidelity methods. For the analyses in this study, the baseline geometry consists of a $1\text{ m} \times 3.06\text{ m}$ rectangular half-wing discretized by 50 spanwise panels and mirrored across the symmetry plane. The thickness-to-chord ratio and location of maximum thickness from the NACA-0012 airfoil are used in the computation of skin friction drag. The geometry is parametrized using B-splines to interpolate variable changes to the geometry. Using B-splines allows a reduction in the number of design variables so that 50 panels can be manipulated with only 9 spanwise control points.

3.1.5 Optimization convergence

All of the optimization results presented herein are converged to a feasibility tolerance of at least 1×10^{-6} , which corresponds to six digits of accuracy in the lift coefficient. The optimality tolerance for every result is at least 2.5×10^{-4} , and for the majority of the results, it is less than 1×10^{-5} .

3.1.6 Optimization Problem

The complete optimization problem of ADODG Case 6 (with some minor alterations) is defined in Table 3.2. Twist variables are defined at every axial control point except at the root and α is used to match the C_L constraint. The chord variables affect all nine axial control points. In the Euler cases, when local shape variables are inactive, the chord scales while maintaining constant t/c . Dihedral is defined as the vertical displacement of each axial control point. Additionally, the FFD section corresponding to each axial control point rotates to align the twist rotation and any shape variable displacements to be perpendicular to the wing surface. Sweep is defined as the streamwise displacement of each axial control point. In the official ADODG case, the limits on sweep extend to 1 m forward and backward in the streamwise direction. In our optimizations, we reduced these limits to 0.5 m to avoid problems with excessive mesh warping. Both dihedral and sweep are fixed at the root. As explained previously, the local shape variables perturb the wing cross-section perpendicular to the wing surface, such that they have some dependence on the dihedral variable. All control points in a given section are perturbed in a uniform direction, as indicated in Figure 3.2. The planform area, S , is computed as the area of the wing projected onto the x - z plane. Constraints for volume (V) and thickness (t) are handled by first setting up a 2D grid of points inside the surface of the wing. Then these points are projected to the surface of the wing to create a 3D grid confined within the wing. We compute V as the sum of the cell volumes and t as the difference between the projected points on the upper and lower surface. The thickness constraints are evaluated at ten uniformly spaced chordwise locations ranging from $0.005c$ to $0.99c$ for ten sections along the span. There are an additional eight thickness constraints added in the wingtip cap, making a total of 108 thickness constraints. We removed the root bending moment constraint from the official optimization problem. The root bending moment constraint is related to structural requirements, and in these results we want to consider solely the influence of aerodynamics on the wing design. The optimization cases treated in Section 3.2 are subproblems of this full problem, and the variables and constraints are defined as stipulated in this full problem description unless otherwise stated. All optimizations are run at Mach 0.5 and a Reynolds number of 5 million.

Table 3.2: ADODG Case 6 Optimization Problem Statement

		Quantity	Lower	Upper	Units
minimize with respect to	C_D	1	–	–	–
	α	1	–3.0	6.0	degrees
	γ	8	–3.12	3.12	degrees
	chord	9	0.45	1.55	m
	dihedral	8	–0.45	0.45	m
	span	1	2.46	3.67	m
	sweep	8	–0.5	0.5	m
	shape	108	–0.5	0.5	m
	Total	143			
subject to	C_L	1	0.2625	0.2625	–
	S	1	3.06	3.06	m ²
	V	1	V_0	–	m ³
	t	108	$0.5t_0$	$1.5t_0$	m
	Total	112			

3.2 Parametrization Studies

3.2.1 Twist Optimization

We begin with a simple twist optimization problem as a means of verification for our optimization framework. The twist optimization case has a long theoretical history and has also been extensively studied as a numerical optimization problem in ADODG Case 3^{86–88}. The elliptical twist distribution is well known as the theoretical optimum for this case because it generates the constant spanwise downwash required to minimize induced drag⁸⁹. This theoretical result is a useful metric with which to gauge the performance of our optimization framework. Figure 3.4 shows the optimal twist and lift distributions for the three levels of fidelity. We see very similar trends from each of the analyses, but there is a noticeable offset between the lift distributions from ADflow and the VLM results. This discrepancy is due to the rounded wingtip cap used for the Euler and RANS geometries. In the lifting-line model, the entire span is used to generate lift, whereas with the wingtip cap, the leading and trailing edges are truncated at 3 m of span and the last 0.06 m of span is incapable of generating the lift required to complete the elliptical distribution. As a result, the optimizer converges to a wing that generates an elliptical lift distribution extending from the root to the edge of the wingtip cap, represented by the shifted elliptical curve in Figure 3.4. The VLM results match the original elliptical curve because they are obtained from a 2D surface for which the problems discussed above are irrelevant. Table 3.3 lists the drag counts of the optimized wings and the percent difference from the baseline drag value, $\% \Delta C_D$. As an added verification, we experimented with varying the number and spacing of the twist variables along the span and also started the optimization from ten random starting points. All of

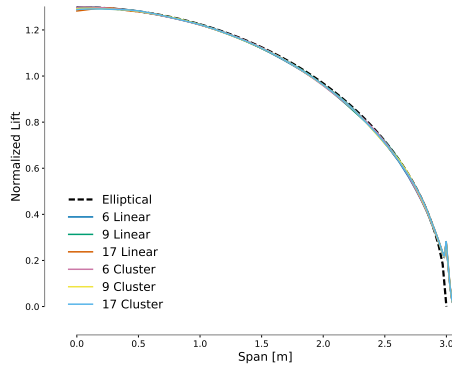


Figure 3.3: Elliptical distribution can be reached with different numbers of FFD control sections and distributions.

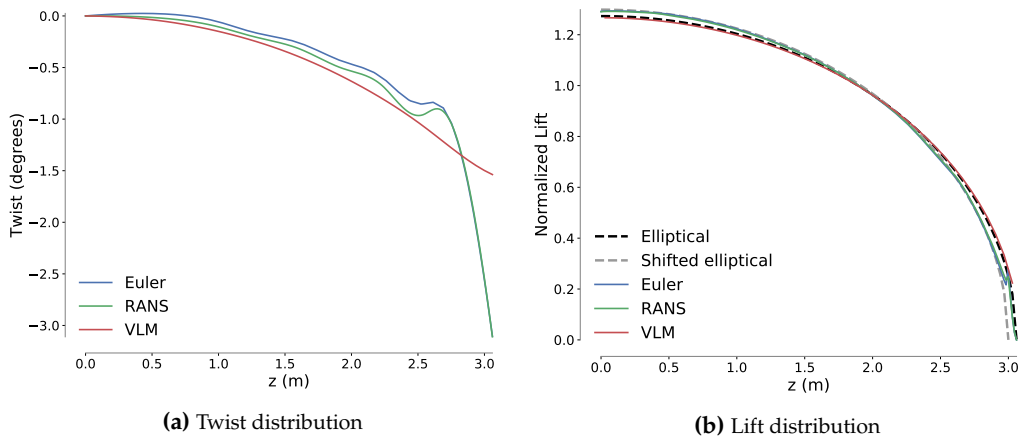


Figure 3.4: Wings optimized with respect to twist variables generate an elliptical lift distribution.

Table 3.3: Twist optimization results

Case	Grid	C_L	Drag counts	$\% \Delta C_D$
Euler	L3	0.2625	42.105	-1.38
RANS	L2	0.2625	143.377	-0.44
VLM	-	0.2625	128.779	-0.42

these variations yielded consistent results, which leads us to confirm the theoretical assertion that there is a single twist distribution that produces the lowest drag.

3.2.2 Chord Optimization

Different behaviors arise when drag is minimized by varying the chord distribution while keeping the twist constant. Theoretically, the elliptical chord distribution should be optimal for minimizing induced drag, and historically, this concept has been put to the test in the design of actual aircraft, most notably the Supermarine Spitfire. The optimization problem is to minimize drag with respect to the chord distribution, subject to the constraint that $C_L = 0.2625$. Since the lift coefficient

is normalized by S , which may vary with changes in the chord, we must also constrain S to be constant. The inviscid and viscous results are henceforth discussed separately for each case to better explore and highlight the unique characteristics of each.

Euler

For the Euler chord optimizations we use the L3 mesh and change the bounds on the chord variables to (0.1, 2.0) to allow the planform to match an ellipse as close as possible. As in the twist optimization, the inviscid chord optimization yields predictable results. We take this predictability as an opportunity to test the sensitivity of the result to the chosen parameterization. We vary the number of spanwise FFD sections and corresponding chord variables and also compare the difference between scaling about the trailing edge and the quarter-chord. As shown in Figure 3.5, regardless of these modifications, each optimization converges to an elliptical lift distribution. The effect of varying the number of FFD control sections is marginal. However, we can see from the results in Table 3.4 that scaling the chord about the trailing edge allows more than double the drag reduction compared with scaling about the quarter-chord. This discrepancy was previously observed in a thorough investigation of the differences in induced drag performance between elliptical and crescent (elliptical with straight trailing edge) planforms by Smith and Kroo⁹⁰. Theoretically, both of these planforms should produce a constant downwash and generate minimum induced drag. Smith and Kroo point out that these optima were derived using lifting-line theory, which relies on a tip vorticity of infinite strength to complete the constant downwash distribution. In nature, such a singularity is impossible, and therefore, even for a perfectly elliptical wing, there is a slight upwash at the tip. They found that the crescent wing produced a more constant downwash near the tip, which led to a more nearly elliptical span loading and a lower overall induced drag than the elliptical wing. This finding may seem to contradict Munk's well-known Stagger Theorem⁸⁹, which maintains that any streamwise arrangement of lifting elements should generate the same induced drag. However, the Stagger Theorem is also based on lifting-line theory, and suffers from the same assumptions in the translation to reality. Thus, although theoretically, a straight wing and a swept wing of the same span and chord distribution generate equal induced drag, in reality, differences in tip downwash will most likely cause slight differences in induced drag too.

To test the multimodality of this problem, we started the optimization from ten random starting points. Nine of the starting points converged to the elliptical planform and one failed prematurely due to mesh warping errors. These results indicate that there is no multimodality in chord optimization for inviscid flow.

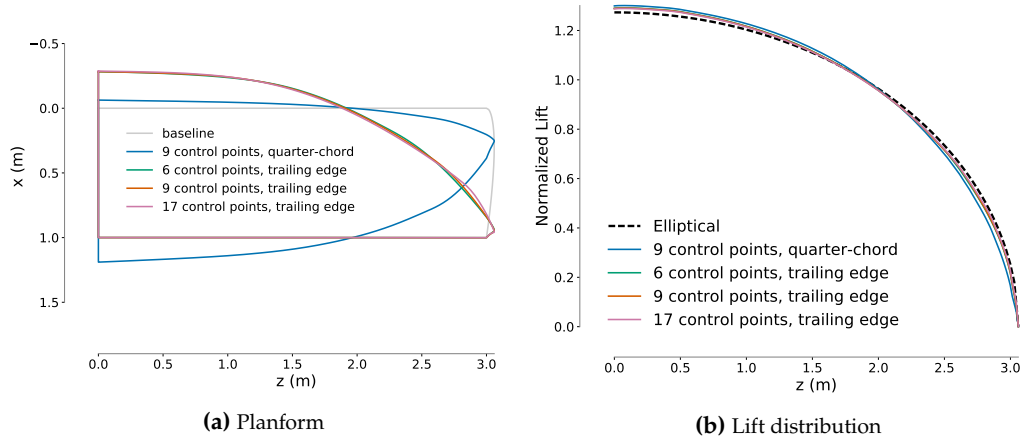


Figure 3.5: Elliptical planform and lift distributions obtained from optimization with respect to chord variables.

Table 3.4: Chord optimization results

Case	Grid	C_L	Drag Counts	$\% \Delta C_D$
Euler 1/4 chord	L3	0.2625	41.726	-2.27
Euler 6 trailing edge	L3	0.2625	40.439	-5.28
Euler 9 trailing edge	L3	0.2625	40.412	-5.35
Euler 17 trailing edge	L3	0.2625	40.356	-5.48
RANS	L2	0.2625	142.544	-1.02
RANS mode 1	L1	0.2625	129.214	-1.02
RANS mode 2	L1	0.2625	129.178	-1.05
RANS monotonic	L1	0.2625	129.296	-0.96
RANS	L1	0.5	234.663	-1.23
RANS	L1	0.8	503.049	-3.85

RANS

When adding viscous effects to the chord optimization, it is important to consider the tradeoff between induced drag and parasite drag, and its relationship to chord length. For the purposes of this discussion, since we are dealing with subsonic flow of a fairly streamlined geometry, we can reasonably assume that the majority of the viscous contribution to drag is made up of skin friction drag. However we do recognize that viscosity also introduces pressure drag due to separation, especially as C_L increases. While induced drag is sensitive to the spanwise distribution of lift, skin friction drag is highly dependent on the local chord length. The shear stress at the wall is directly related to the velocity gradient normal to the wall. For 2D laminar flow over a flat plate,

$$\tau_w = \mu \frac{du}{dy} \quad (3.1)$$

This equation provides an approximation to the shear stress on an airfoil. At the leading edge, the boundary layer is very small and the velocity changes rapidly over a small distance, resulting in a large shear stress. As x/c increases, the boundary




Figure 3.6: The viscous drag on a flat plate in laminar flow can be reduced by increasing the average BL thickness over the plate. ($Re = 10^6$)

layer (BL) fills out and the velocity gradient at the wall becomes much more mild. The result is that extending the chord reduces the average drag per unit length. This relationship is well known and Blasius⁹¹ provided the following analytic solution for the flat plate case,

$$c_d = \frac{1.328}{\sqrt{Re}}, \quad (3.2)$$

where the Reynolds number is based on the local chord. Multiplying by the local chord length we find that for constant flow conditions, the skin friction drag per unit span is proportional to the square root of the local chord length, i.e.,

$$d \propto \sqrt{c} \quad (3.3)$$

To illustrate this point, imagine we want to minimize the drag of a flat plate in laminar flow at zero angle of attack with a fixed planform area. If the chord distribution and span are variables, the chord distribution will grow to its upper limit and the span will adjust to satisfy the area constraint. However, if span is fixed, or has a hard lower limit, an interesting compromise takes place. The optimal chord distribution will have the maximum possible extent of the span at the upper limit of the chord variable. For example, if we split our plate into two independent sections and incrementally add Δc to the chord of one section while subtracting the same Δc from the other, we get a decrease in skin friction drag as shown in Figure 3.6. In the absence of other constraints, lower skin friction drag can always be achieved by transplanting wing area from a thin-BL region to a thick-BL region. The spanwise location of the maximum chord region is irrelevant, and as such, a purely skin friction drag minimization problem theoretically has an infinite number of local minima.

The problem becomes more complex when the objective function is a combination of both skin friction and induced drag. Minimum skin friction drag favors

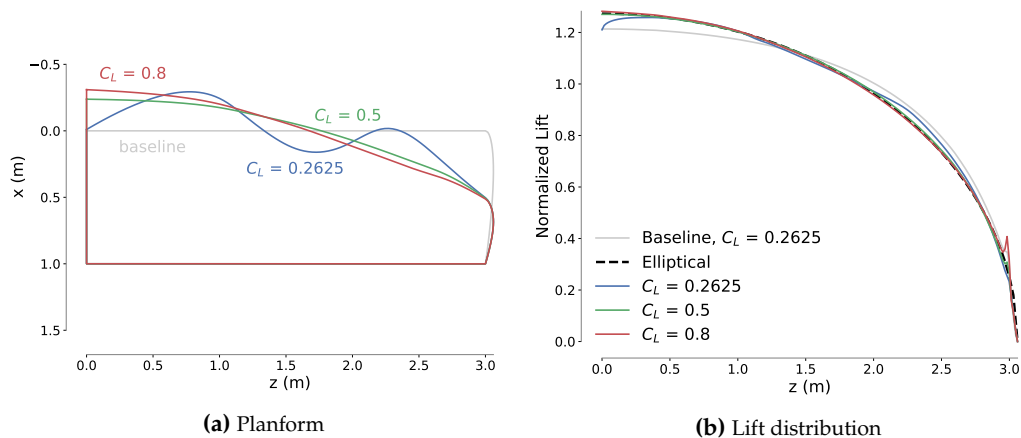


Figure 3.7: The optimal chord distribution for viscous flow is heavily dependent on C_L .

radical changes in the chord distribution to maximize thick-BL coverage, but optimal inviscid drag calls for an elliptically tapered wing. The optimal planform shape will balance these considerations taking into account the relative weights of each drag component. For instance, if drag is mostly induced (i.e., at high C_L), the planform will take on a nearly elliptical profile. However, if we make some slight modifications to the planform, we may get some improvement in the skin friction drag while not straying too far from the elliptical lift distribution. This line of reasoning helps to explain the results we get from the RANS chord optimization, shown in Figure 3.7. These results are obtained using the RANS L1 mesh. For a given C_L , we can estimate the minimum possible contribution of induced drag to the total drag using Prandtl’s induced drag approximation $C_{D,i} = C_L/\pi AR$. For $C_L = 0.2625$, parasite drag makes up roughly 70% of the total drag, and the optimal planform shape is far from elliptical. As C_L increases and induced drag increases relative to parasite drag, the optimal design more closely approaches the elliptical planform. However, even when C_L is relatively low, the optimal lift distribution oscillates close to the elliptical profile. This is not surprising because in reality, even a rectangular wing has a span loading not far removed from the elliptical ideal. Figure 3.7 shows the lift distribution of the baseline wing at $C_L = 0.2625$. As seen in the Euler results in Table 3.4, going from a rectangular wing to an elliptical wing produces a 2–6% drag reduction. So, while the elliptical planform is ideal, the rectangular wing is certainly not a terrible starting point. This accounts for the rather meager gains in aerodynamic performance observed in the twist and chord optimizations overall.

To address whether this tradeoff between parasite and induced drag leads to multimodality at low C_L values, we first turn to a simpler aerodynamics model. Figure 3.8 shows the results of optimizing 20 random initial planforms at five dif-

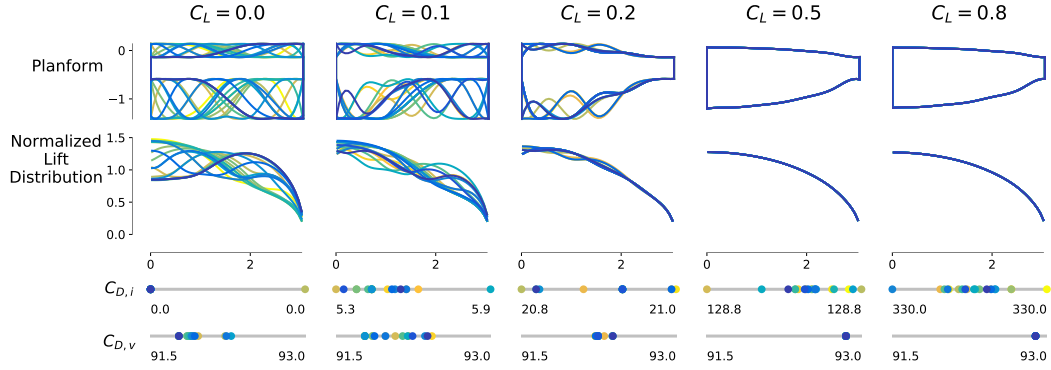


Figure 3.8: Optimization using OpenAeroStruct confirms multiple optimal chord distributions at low C_L .

ferent C_L values with OpenAeroStruct. A tradeoff exists between minimizing skin friction drag by maximizing chord and forming an elliptical planform to minimize induced drag. As expected, based on our previous discussion, the deviations from the elliptical lift distribution increase dramatically as C_L drops to zero. A corresponding decrease in skin friction drag is small but noticeable. As C_L increases, induced drag makes up a greater portion of the total drag, and thus the planform approaches an elliptical shape and multimodality decreases. In the high-fidelity case we see similar trends. We optimize three randomly generated planforms at $C_L = 0.2625$ with the RANS L1 mesh and the optimizer converges to two different planform modes. Wary of this result, we vary the number of control points used as design variables to rule out the possibility that our parametrization is biasing the results. We optimize the baseline geometry with 6 and 17 spanwise variables and compare the result with the baseline optimization with 9 spanwise variables. Despite the variance in flexibility allowed by the number of spanwise control points, the optimizer converges on a very similar, albeit not identical, planform shape in all three cases. These experiments suggest that there are a limited number of modes with which the optimizer can minimize skin friction drag while still maintaining a sufficiently elliptical lift distribution so that the total drag decreases. The results for both of these tests are shown in Figure 3.9. Incidentally, we began the chord optimization tests with the RANS L2 mesh, but found a lack of multimodality in the low C_L cases. Since this did not agree with our hypothesis, we refined the mesh and found that the finer mesh yielded sharper spanwise curvature due to the refinement of the spanwise grid spacing. We surmise that the coarseness of the L2 mesh causes an increase in drag when the optimizer attempts to create the large-amplitude spanwise variations seen in Figure 3.9 and thus artificially limits the design space. However, it should be noted that the same physical phenomenon is apparent in both meshes, although its effect is dampened due to the coarseness of the mesh.

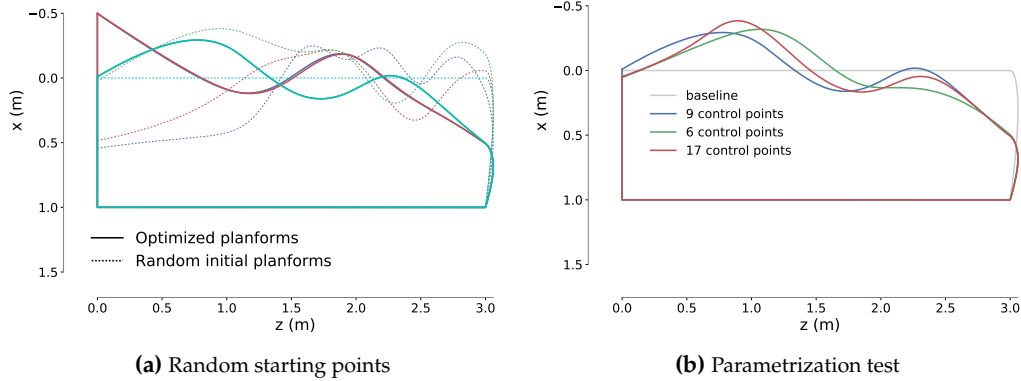


Figure 3.9: (a) Three optimizations from random chord distributions yield two local minima. (b) Varying the number of chord variables changes the optimal distribution slightly.

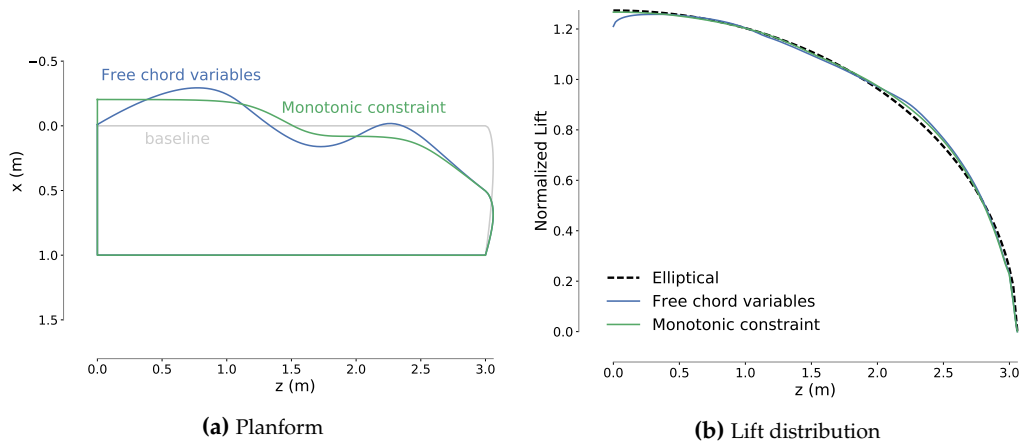


Figure 3.10: Forcing the chord distribution to decrease monotonically increases drag by less than one count.

While these results are intriguing, we are ultimately interested to see how much of a benefit these spanwise oscillations really provide. The VLM results suggest a drag decrease on the order of one to two counts. We add a monotonic constraint to the optimization of the L1 grid as a crude substitute for a more meaningful constraint, such as manufacturing cost and structural considerations. This constraint forces the chord distribution to decrease monotonically from root to tip. The result is plotted in Figure 3.10 and detailed in Table 3.4. Adding the monotonic constraint only increases the drag by a fraction of a drag count. This difference is not meaningful because it is within the modeling error. Furthermore, even if this difference were meaningful, it would be hard to justify the added manufacturing costs and structural penalties that would accompany building a wing with such spanwise curvature.

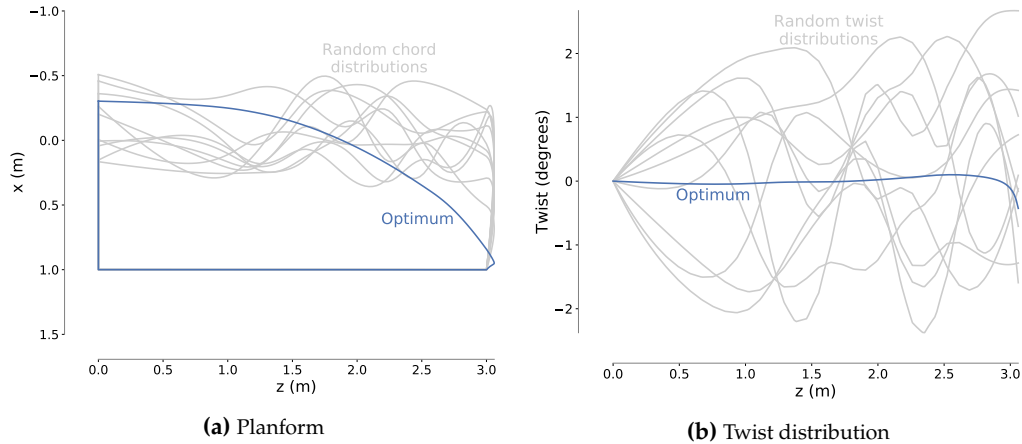


Figure 3.11: Optimizing from 10 unique starting points with respect to chord and twist variables using Euler analysis yields a single optimum.

3.2.3 Chord and Twist Optimization

Since we can reach an elliptical distribution with only chord design variables or only twist design variables, one may assume that a combination of both twist and chord variables would be redundant and lead to multimodality. We explore this possibility in this section.

Euler

For the chord and twist optimization with Euler analysis we use the Euler L3 mesh and nine axial control points, corresponding to nine chord variables and eight twist variables. Once again, the limits on the chord variables are relaxed to $0.1 \leq c \leq 2.0$ for the Euler cases. When initialized with ten different geometries, each generated with random chord and twist distributions, the optimizer converges to a single optimal solution. The optimal chord and twist distributions are superimposed over the randomly generated seeds in Figure 3.11. The uniformity of these results suggest that the chord and twist variables are not necessarily interchangeable. In Table 3.5, we see that the addition of twist variables makes a negligible improvement in drag compared with the chord-only optimization (Table 3.4).

RANS

The RANS optimizations are run using the RANS L1 mesh and the nominal 9-section FFD. The results are displayed in Figure 3.12. We expect to see the oscillatory behavior observed at low C_L in the chord optimization to be exaggerated in this case because the twist can compensate for the deviations from the elliptical lift distribution that oscillations in the chord distribution would otherwise cause. The

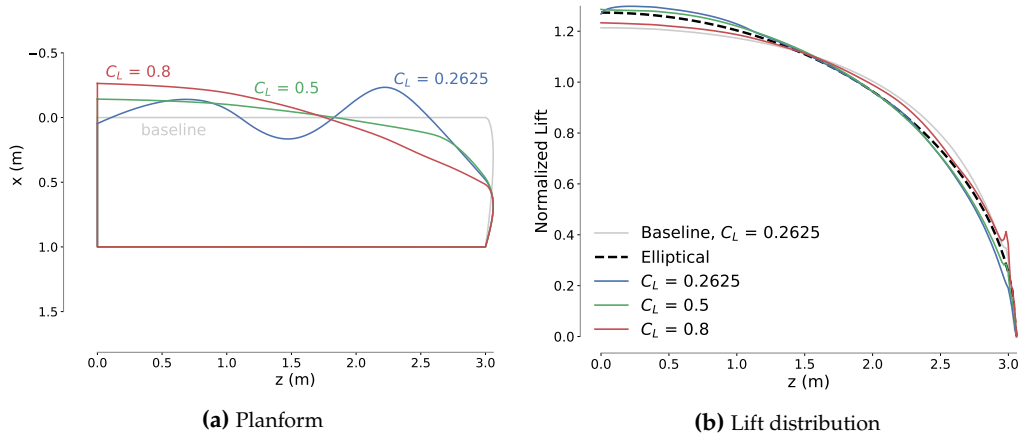


Figure 3.12: An optimization with both twist and chord variables tolerates more variation in the chord distribution at a lower C_L . As C_L increases, the planform shape oscillations disappear.

Table 3.5: Chord and twist optimization results

Case	Grid	C_L	Drag Counts	$\% \Delta C_D$
Euler 1/4 chord	L3	0.2625	41.517	-2.76
Euler	L3	0.2625	40.408	-5.35
RANS	L1	0.2625	128.802	-1.34
RANS	L1	0.5	233.591	-1.68
RANS	L1	0.8	480.496	-8.16

higher C_L results, like the Euler results, show little evidence of multimodality due to redundancy in the variables. For $C_L = 0.2625$, the oscillations in the chord distribution are more prominent, which implies that the addition of twist variables grants more freedom to the chord variables to minimize skin friction drag. Once again, the improvement in drag gained from the chord and twist optimization is fairly insignificant due to the fact that there is only so much that can be improved with twist and chord variables.

3.2.4 Dihedral and Twist Optimization

Nonplanar wings have the potential to reduce induced drag beyond what is attainable with a planar wing and an elliptical lift distribution. Previous optimization studies have verified this result^{55,92}. We include twist variables in this subproblem because we want to allow the optimizer to converge to an elliptically loaded planar wing if that is the optimal design. However, we first assess the parametrization of the dihedral variables by conducting an optimization with just dihedral. Since we expect the most variation in dihedral to occur toward the wingtip, we want to make sure that our parametrization allows enough flexibility to capture the optimal shape. We vary the number of control points and experiment with spacing them uniformly along the span and clustering them more heavily toward the wingtip. In

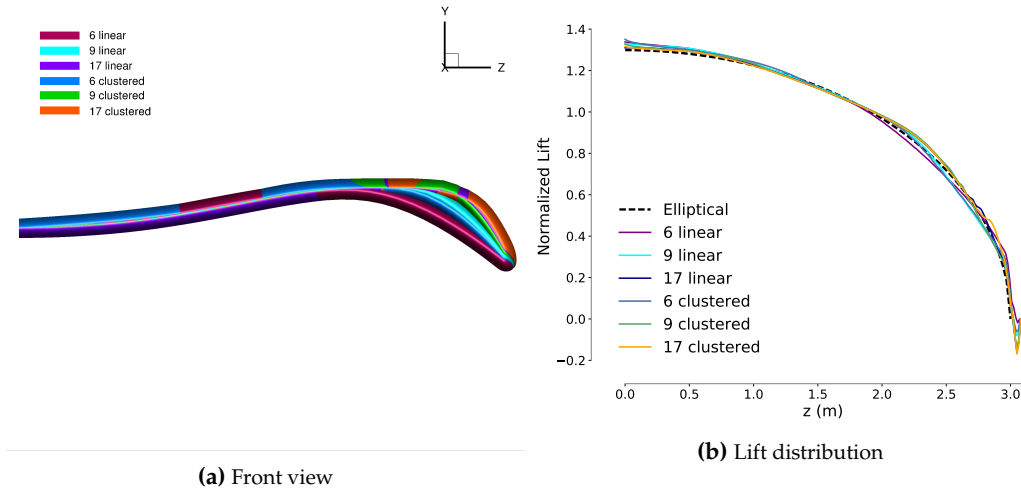


Figure 3.13: Comparison of different quantities and spanwise distributions of control points for an optimization with dihedral variables.

Figure 3.13 we compare the winglet-down optimized results for six different possible parametrizations. All six parametrizations achieved an elliptical lift distribution, but the cases with only six variables prevented the optimizer from converging to the optimal winglet cant angle. We choose to use the FFD with nine control points clustered toward the wingtip as the nominal FFD because it gives a sufficient degree of freedom to the optimizer for the kind of design space exploration we seek to do. Note that we recover two local minima in the dihedral-only optimizations: an upturned winglet and a downturned winglet. The upturned winglet shapes are achieved when starting the optimization from random starting points; Figure 3.13 shows only the optima found when starting from the baseline geometry.

Euler

Now we combine twist and dihedral to see if the coupling between these two sets of variables affects the number of optimal designs. We start with running 15 optimizations starting from unique dihedral distributions. These initial distributions can be seen in the top frame of Figure 3.14. In the optimization with L3 meshes, the optimizer converges on multiple local minima. Five of the optimizations converge to an upturned winglet shape and the rest converge to a downturned winglet. Those that converge to the winglet-down shape share a very similar wingtip design, but vary in the vertical displacement of the main wing. These differences can be seen in the middle frame of Figure 3.14. The next set of optimizations starts from the colored distributions in the middle frame of Figure 3.14, but uses L2 meshes. The designs obtained in the optimization with finer meshes are shown in the final frame of Figure 3.14. The transition to finer meshes reveals that some of the multimodality

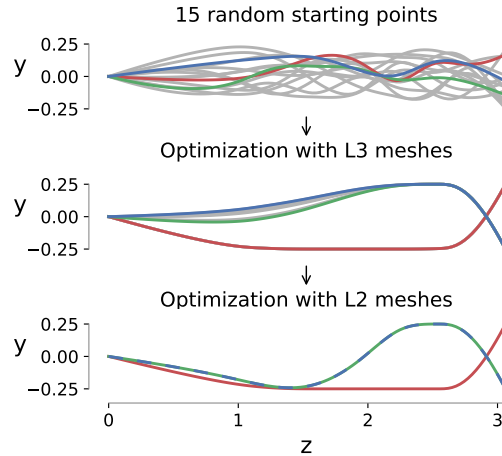


Figure 3.14: Results of optimizing the dihedral distribution first with L3 meshes and then with L2 meshes.

Table 3.6: Dihedral and twist optimization results. The $C_L = 0.8$ result was obtained by regenerating the volume mesh to fix problems with the formation of negative volumes.

Case	Grid	C_L	Drag Counts	$\% \Delta C_D$
Euler	L3	0.2625	39.909	-6.52
RANS	L2	0.2625	143.327	-0.48
RANS	L2	0.5	244.969	-2.79
RANS	L2	0.8	455.178	-15.21

observed in the first optimization was due to the coarseness of the mesh, and not to any physical phenomenon. In the end, the optimizer converged to a wing with an upturned winglet and a wing with a downturned winglet. In inviscid flow, increasing the horizontal and vertical extent of the wing helps to decrease drag. In the downturned winglet design, the optimizer has maximized the utilization of vertical space by changing curvature in the middle of the wing. Although this is optimal for inviscid analysis, the penalty in parasite drag due to the increased surface area would make this wing suboptimal in realistic analysis.

RANS

The addition of viscosity to the model activates a tradeoff between the induced drag improvement from winglet formation and the rise in parasite drag due to an increase in wetted area. Before running the RANS analyses, we explore the design space using OpenAeroStruct to better understand the implications of this tradeoff. Figure 3.15 compares the optimization results for three general cases: inviscid analysis, viscous analysis, and viscous analysis with a monotonic constraint. The monotonic constraint forces the optimizer to form a downturned winglet, but the results are of interest because they give a metric with which to compare the other results. The key in the upper left corner provides labels for the different colors. The

three cases are optimized with 5 different starting points, displayed in the left-most column, and 5 different lift conditions ranging from $C_L = 0.1$ to $C_L = 0.8$. The front views of each optimized wing profile are shown in each cell of the grid. Below and to the right of each wing profile, slider bars show the relative drag differences between the viscous-optimized results and the inviscid-optimized result. The horizontal bar shows the percent difference in inviscid drag, while the vertical bar indicates the percent difference in skin friction drag compared with a viscous analysis of the inviscid-optimized result. Moving from the left to the right of the grid, the trend is an increase in nonplanarity for the viscous results. At low C_L values, when the skin friction drag dominates, the optimizer has an incentive to reduce the arc length of the front view and thus reduce the effective wetted area. Since only small changes to the dihedral variables are possible without increasing the wetted area, the optimizer converges to the same optimum regardless of the starting point. As C_L increases, this incentive diminishes and the optimizer tends toward more nonplanar wing shapes where the induced drag can be minimized. The increased freedom to vary the wing dihedral opens up the design space to multiple local minima. The inviscid optimization converges to three local minima and the viscous optimization for $C_L = 0.8$ converges to two optima: an upturned and a downturned winglet. The downturned winglet appears to offer better drag performance compared with the upturned winglet.

Taking the information from OpenAeroStruct as a reference, we optimize the RANS L2 baseline planform with respect to eight dihedral and twist variables for $C_L = 0.2625, 0.5$, and 0.8 . The results once again exhibit a strong dependence on C_L . As seen in left-hand side of Figure 3.16, the low C_L case deflects only slightly from the baseline geometry. As C_L increases, the vertical extent of the wing expands. Interestingly, when chord variables are added into the optimization, the trends seen in the separate chord and dihedral optimizations seem to be linearly combined with no recognizable coupling between dihedral and chord. This combined dihedral, twist, and chord optimization is presented on the right-hand side of Figure 3.16. In both of these cases, the $C_L = 0.5$ result converges to a down-turned winglet, while the $C_L = 0.8$ result converges to an upturned winglet. We run a dihedral and twist optimization for three wings with randomly distributed dihedral and twist at $C_L = 0.5$ to investigate the possibility of multiple local minima (Figure 3.17). Surprisingly, the optimizer converges to an upturned winglet for all three of these randomly generated wings, despite the fact that a comparison of drag values reveals a preference for the downturned winglet. We conclude that in an optimization where the parametrization allows for winglet formation, a gradient-based optimizer can converge to either an upturned or downturned winglet, depending on the starting position. If removal of one of these optima was desired, it would be trivial




Figure 3.15: We use OpenAeroStruct to quickly explore the design space of the dihedral and twist optimization.

to apply a constraint which would force the winglet in the preferred direction.

3.2.5 Adding span and sweep variables

As a final step before considering the full case, we investigate the effects of adding span and sweep variables to the RANS optimization. Figure 3.18 shows the results of two distinct optimization problems. The left-hand side shows a comparison between optimizing with respect to twist, chord, dihedral, and span at $C_L = 0.2625$ and $C_L = 0.5$. The only geometric constraint is planform area. For $C_L = 0.2625$, the optimizer finds greater benefit in maximizing the chord and does not increase span to reach the upper bound. The span does increase from the baseline case though, and the optimizer adds a slight anhedral to the wing. On the other hand,

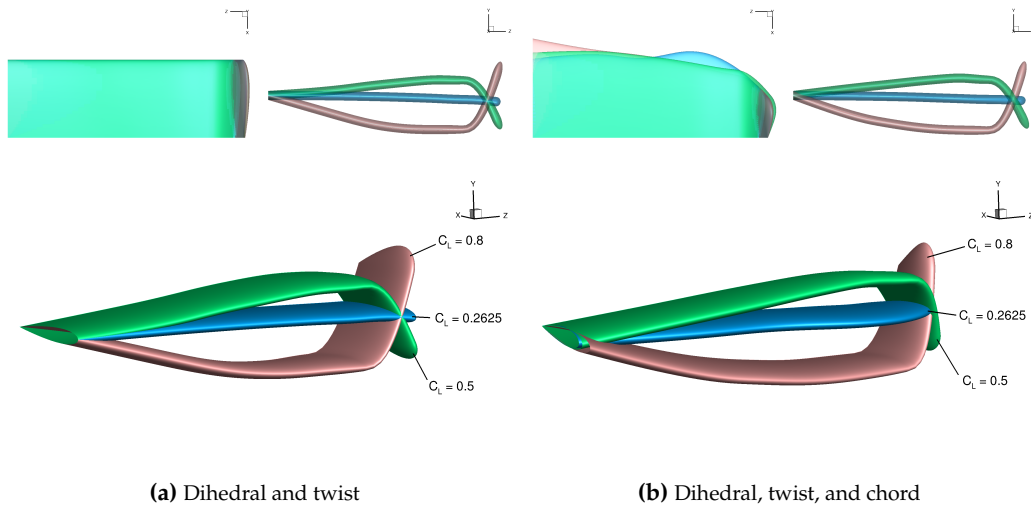


Figure 3.16: (a) A nonplanar wing forms as C_L increases. (b) The chord varies at low C_L and the winglet forms at a high C_L .

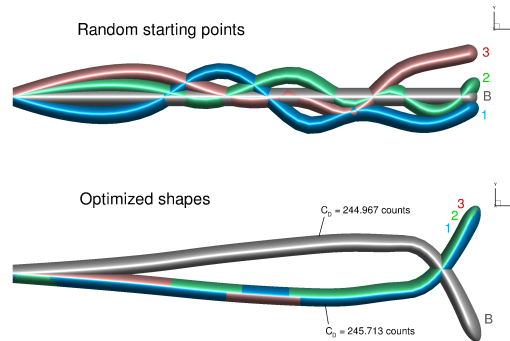


Figure 3.17: The optimizer converges on two local minima for $C_L = 0.5$.

for $C_L = 0.5$, the span reaches the upper bound and an upturned winglet is formed.

When sweep is added to the optimization, we find the optimizer tends to sweep the wingtip back sharply, creating a raked wingtip. When the full bounds recommended in the ADODG case description are used, this sharply swept wing tends to cause mesh warping errors, resulting in negative volumes. To avoid this problem, we reduced the bounds to allow sweep to vary from -0.25 m to 0.25 m. The right-hand side of Figure 3.18 shows the results of this optimization. The most noticeable feature of these results is the formation of the raked wingtip. Once again we notice that the optimizer does not extend to the full bounds of the span variable for the lower lift case. This time, for $C_L = 0.5$, the optimizer converges on a downturned winglet. The appearance of the same swept-back wingtip in both local minima, despite the differences in chord, dihedral, and sweep distribution on the main part of the wing, suggests that the sweep at the wingtip is highly beneficial to drag minimization. For a wing at incidence to the freestream tip sweep can function similarly to dihedral, dispersing the shed vorticity over a larger region and reducing vortex

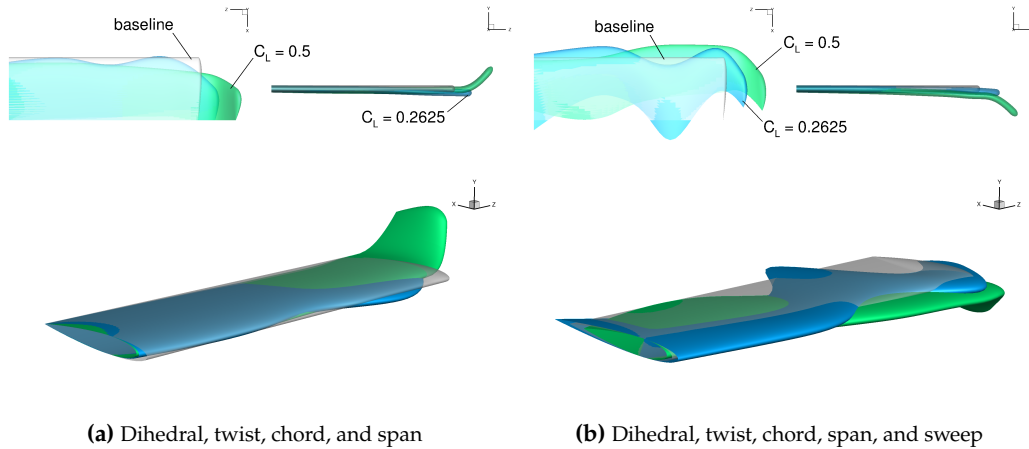


Figure 3.18: RANS optimization results starting from the baseline rectangular wing at different C_L values.

drag. Additionally, the wing could be benefiting from more constant downwash due to the tip sweep, as explained in Section 3.2.2 regarding the crescent wing.

3.3 Full Case Optimization

The results of Section 3.2 emphasized the importance of considering viscous effects when spanwise chord and dihedral variables are included in the optimization problem. Therefore, although the ADODG case calls for Euler analysis, we only report optimization results obtained using RANS analysis in this section. Euler analysis likely would produce multiple local minima, however, our purpose is to study physically meaningful multimodality, and thus we restrict ourselves to RANS analysis. All of these results are produced using the RANS L2 mesh.

The first result presented here is a solution of the optimization problem in Table 3.2. We run the optimization starting from the baseline geometry and two randomly perturbed starting positions. As seen in Figure 3.19, the optimizer converges on a different local minimum in each case. Despite the apparent differences in planform shape, the final drag values are remarkably similar (see Case 1 in Table 3.7). Following the trend of these results, we suppose that starting the optimization from additional points in the design space would reveal more local minima. In keeping with our purpose, we do not attempt such an investigation, but we would like to point out some of the trends that have reappeared from Section 3.2 in these results. True to our previous hypothesis, the chord variation is multimodal—even more so than before because for a given chord distribution the optimizer has more design variables with which to tailor the wing shape to minimize induced drag. Each optimal design has a unique sweep distribution, but there is a commonality of sharp sweep-back at the wingtip. As expected from our study of the dihedral

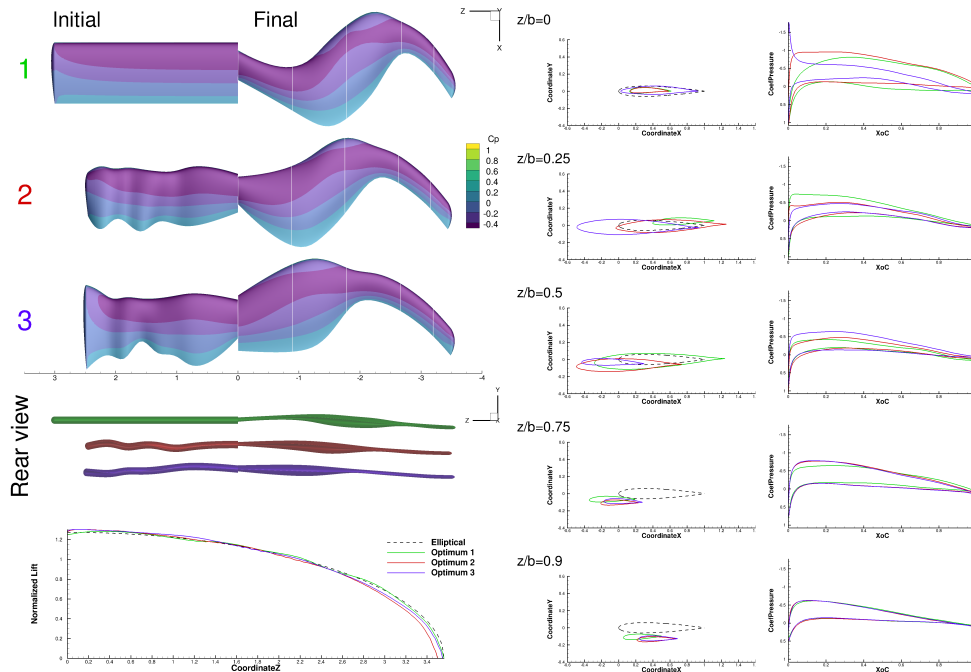


Figure 3.19: Three local minima found with all variables active.

variables, there is very little y -displacement in any of the three solutions due to the low lift coefficient and the relative unimportance of minimizing induced drag. Instead of pushing the span to its maximum allowable extent, the optimizer uses the available planform area to maximize the chord and reduce skin friction drag. In all three cases, the span increases to a different value, implying that the optimum aspect ratio varies depending on the shape of the wing. In terms of drag reduction, all three solutions improve by nearly 11%. The breakdown between viscous and pressure drag in Table 3.7 reveals a significant reduction in pressure drag. Up to this point in our investigation we have not used shape variables to tailor the airfoil cross-section of the wing. It appears that the optimizer has taken advantage of the shape variables to streamline the wing shape and reduce drag due to viscous separation. This can also be seen in the pressure plots on the right of Figure 3.19, where the flattened profiles differ significantly from the well-known steep pressure gradient of the baseline NACA-0012 airfoils.

In the remaining results, we consider different methods of constraining the problem to manage the multimodality. We first experiment with adding the root bending moment constraint back into the problem. The result is a planform similar to Optimum 1 in Figure 3.19, only with a shorter span to satisfy the added constraint and a corresponding increase in drag (see Case 2 in Table 3.7). Since adding this constraint does not appear to attenuate the multimodality in the chord distribution we do not investigate further. Next we reintroduce the constraint forcing the chord distribution to decrease monotonically from the root to the tip. With this constraint,

Table 3.7: Full optimization problem results. Percentage change is referenced to the baseline wing.

Case	Notes	#	C_L	C_D (cts)	$\% \Delta C_D$	$\% \Delta C_{D,v}$	$\% \Delta C_{D,p}$		
1	All variables	1	0.2625	128.358	-10.87	-1.3	-21.7		
		2	0.2625	128.847	-10.53	-1.0	-21.3		
		3	0.2625	128.888	-10.50	-0.7	-21.6		
2	Root bending moment constraint	1	0.2625	131.903	-8.41	-2.8	-14.7		
		3	Monotonic chord constraint	1	0.2625	128.157	-11.01	-0.8	-22.5
				2	0.2625	128.116	-11.04	-0.3	-23.1
3	Monotonic chord constraint	3	0.2625	128.115	-11.04	-0.3	-23.1		
		4	Linear chord and sweep distributions	1	0.2625	134.011	-6.95	0.6	-15.4
				2	0.2625	134.011	-6.95	0.6	-15.4
3	0.2625			134.010	-6.95	0.6	-15.4		

the optimizer converged on two different local minima in three different runs. The results are shown in Figure 3.20 and tabulated in Table 3.7 (Case 3). For both of these optima, the monotonically decreasing chord distribution resulted in slightly better performance than any of the optima found in the previous optimization. Adding the monotonic constraint gave the optimizer freedom to extend the span to its maximum allowable extent, resulting in better induced drag performance. This can be seen in the slightly better reduction in pressure drag for the Case 3 optima compared with the Case 1 optima. These results corroborate our previous assertion that minimizing skin friction drag with respect to the chord distribution is a multimodal problem. The optima found in Case 3 were present in the design space of Case 1, but the optimizer was unable to find them because of the multimodality due to the chord variables. Importantly, failing to find the optimal chord distribution prevented the optimizer from maximizing span and achieving the optimal aspect ratio. The major difference between the two optima in Case 3 is the sweep distribution. Optimum 1 follows a zig-zag pattern, sweeping back, then forward, then back again at the tip. Optima 2 and 3 have a more gradual sweep distribution, sweeping forward until midspan and then back to the tip. Interestingly, the same sweep patterns are observed in Figure 3.19. The distribution of anhedral is nearly identical across the optimized designs, but it varies only slightly from a flat wing, indicating the lack of benefit to be gained from a winglet at low C_L .

While these optimized wings certainly perform better than the baseline wing, there is no disputing the fact that they are unconventional—possibly too unconventional to be practical. The final optimization case attempts to address the tradeoff between performance and practicality. For this final case, we add a constraint that the chord and sweep distribution must vary linearly along the span of the wing. In three different runs, the optimizer converges on a single design, shown in Figure 3.21. This design appears much more conventional, yet it suffers a 4% increase in drag compared with the optimized designs from Cases 1 and 3. The optimum design has maximum sweep and minimum tip chord. The span upper bound is not reached, probably due to the need to maintain a relatively large root chord in order

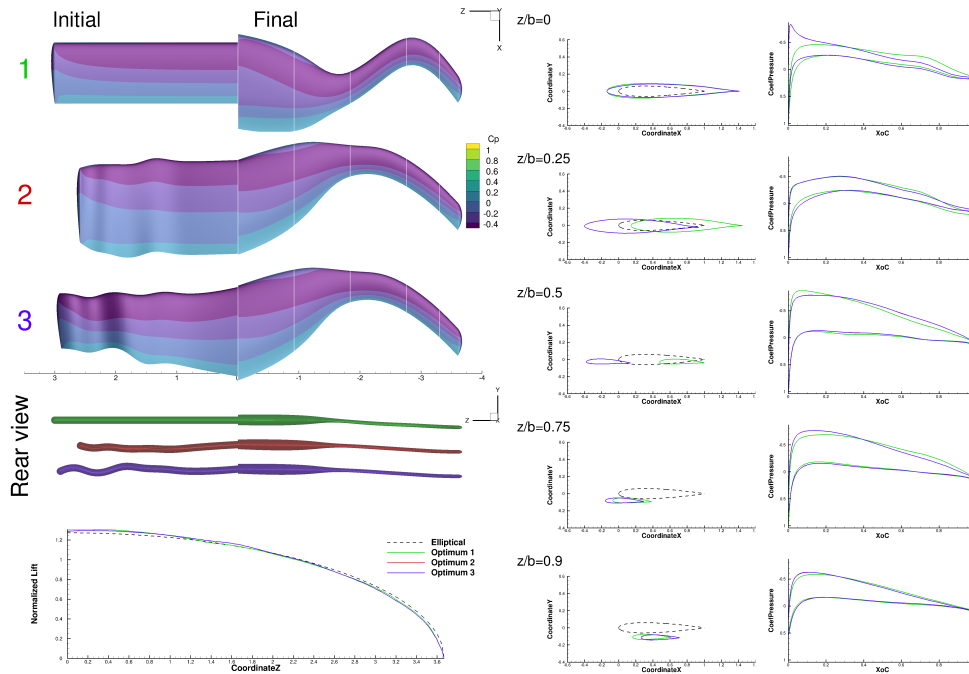


Figure 3.20: Two local minima found when monotonically decreasing chord is enforced.

to not cause a greater increase in skin friction drag.

3.4 Sweep variable study

The distinctive curvature in the sweep distribution of the best-performing optimum in Table 3.7 warrants further analysis. At subsonic speeds, the effect of sweeping the wing should be marginal⁸⁹. Using optimization we can determine whether or not our modeling approach reflects this hypothesis. We performed an additional optimization with an identical problem formulation except that the sweep is fixed to maintain the quarter-chord parallel with the z -axis. The optimized straight wing achieves nearly the same same drag reduction as the scimitar wing. Table 3.8 shows that the difference is less than one drag count. The optimized geometries are depicted in Figure 3.22 along with airfoil and C_p slices at various spanwise locations. The airfoil slices are quite similar despite the differences in planform shape, indicating that there is little coupling between sweep and airfoil shape at this subsonic, low- C_L flow condition. Figure 3.23 shows the variation of aerodynamic and geometric properties along the span of the two wings at the design ($C_L = 0.2625$) and zero-lift conditions. The chord distribution and the lift distributions at the design point are identical. Likewise, the spanwise variation in viscous drag is minimal. The major difference between the two wings is the spanwise distribution of pressure drag. The pressure drag of the straight wing is roughly constant along the span, while that of the scimitar wing has a similar curvature to its sweep

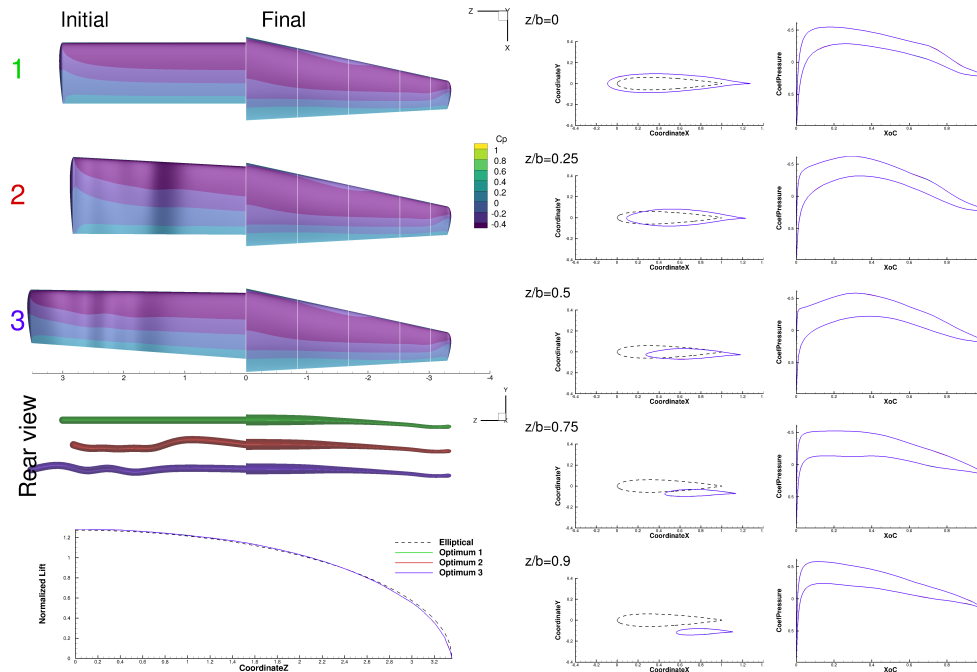


Figure 3.21: One optimal design found when chord and sweep distributions are constrained to be linear.

Table 3.8: Sweep optimization problem results. Percentage change is referenced to the baseline wing.

Case	C_L	C_D (cts)	$\% \Delta C_D$	$\% \Delta C_{D,v}$	$\% \Delta C_{D,p}$
Scimitar	0.2625	128.116	-11.04	-0.3	-23.1
Straight	0.2625	129.036	-10.40	-0.03	-22.2

distribution. Interestingly, the pressure drag of the scimitar wing has a higher peak than the straight wing, but this is offset by small regions of negative drag at the root and the tip. It is unclear whether this effect is an artifact of the numerical modeling or a relevant physical phenomenon. However, given that the integrated effect of the scimitar sweep distribution is only marginally better than the straight wing, we conclude that for this design case and flow condition, the optimum design is relatively insensitive to the sweep variable.

3.5 Summary

Wing aerodynamic optimization with respect to planform variables involves a strong tradeoff between induced and parasite drag. The inclusion of a viscous drag model in the optimization is necessary to fully account for this tradeoff and obtain physically meaningful results. When minimizing drag with respect to chord distribution at low C_L values, this tradeoff causes the optimizer to form nonintuitive wavy chord distributions that minimize skin friction drag at the cost of a perfectly

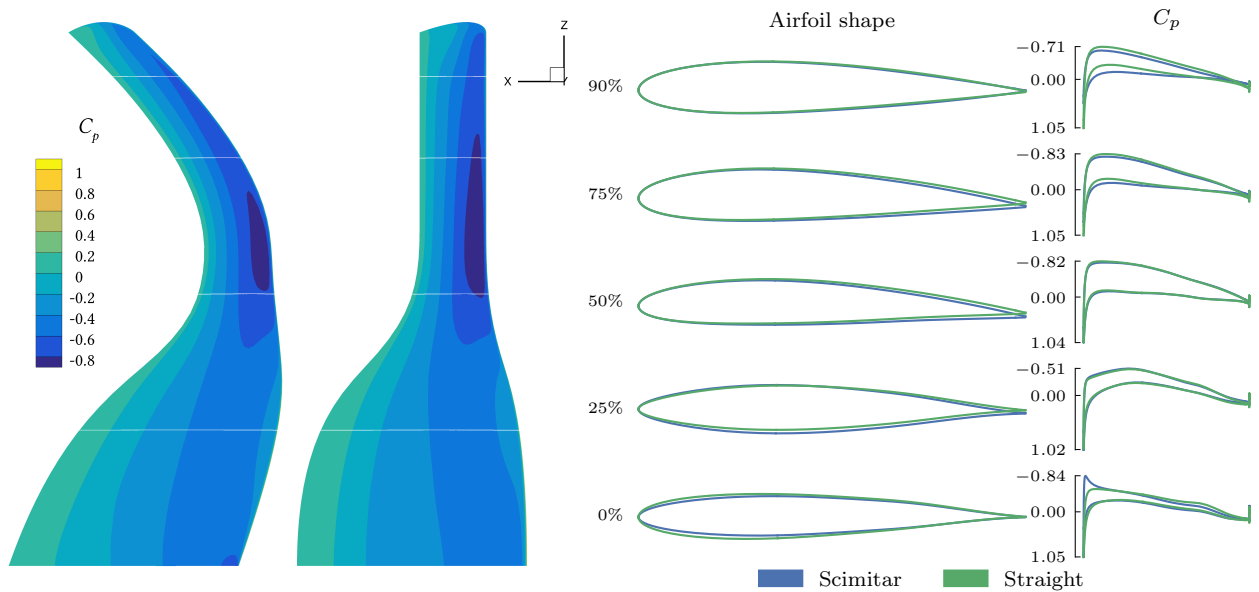


Figure 3.22: Comparison of scimitar and straight wing geometries and C_p distributions.




Figure 3.23: Comparison of spanwise variation on scimitar and straight wings.

elliptical lift distribution. This tradeoff introduces multimodality into the design space, which we explained by analyzing a canonical flat plate with zero lift. In this extreme case, the optimal chord distribution is governed by the motivation to increase the chord as much as possible to reduce the skin friction drag per unit area while meeting the wing area constraint, but multiple chord distributions yield the minimum drag solution. As C_L increases, this multimodality is attenuated by the increasing relative importance of induced drag because the induced drag is also coupled with the chord distribution.

The benefit of nonplanarity is strongly correlated with the drag tradeoff as well. When optimizing with respect to dihedral variables, the optimizer forms a winglet, for both $C_L = 0.5$ and $C_L = 0.8$. The optimizer also displaces the main wing vertically in the opposite direction of the winglet to maximize the vertical extent of the wing and thus the vorticity sheet. This trend is amplified for $C_L = 0.8$ compared with $C_L = 0.5$. On the other hand, for a low C_L condition, the decrease in induced drag due to nonplanarity does not outweigh the increase in parasite drag due to the additional wetted area. Thus, for high C_L conditions, dihedral variables add multimodality to the design space. We find that a down-turned winglet reduces drag more than an upturned winglet, but both are local minima. A wing optimized with respect to both chord and dihedral variables appears to linearly combine the trends found in the separate cases: at low C_L the planform is wavy with indiscernible dihedral, and at high C_L the wing tapers to a winglet.

The optimal aspect ratio also depends on the lift condition. A short, stubby wing is best for low-lift flight dominated by skin-friction drag. On the other hand, when induced drag is the principal drag component, a wing with a high aspect ratio is desirable. Importantly, the multimodality in the chord distribution can prevent the optimizer from achieving the optimal aspect ratio to minimize induced drag. This is seen in the comparison of Figures 3.19 and 3.20, where the restriction of the design space to monotonically decreasing chord distributions actually enables the optimizer to find a chord distribution that allows a higher aspect ratio and lower overall drag. The designs in Figure 3.20 are present in the design space of Figure 3.19, yet the optimizer does not find them due to multimodality.

When the optimizer has freedom to control sweep, it frequently converges to a scimitar-like wing with a sharply swept wingtip. However, when the wing is parametrized to have zero sweep, the drag is only slightly higher than the scimitar wing. This indicates that for subsonic flow the design space of the sweep variable is relatively flat and variations in sweep have minimal impact on the total drag of the wing.

It is crucial to consider viscous effects when the geometric parametrization allows the development of a nonplanar wing. In such cases, reducing induced drag

is an incentive to increase dihedral and form a winglet, which also increases the total surface of the wing. Viscous effects must be considered in order to balance the tradeoff between induced drag and parasite drag regarding the formation of a winglet. All of these results strongly support the use of viscous analysis over inviscid analysis in realistic wing design optimization.

Although we did find evidence of multimodality for various parametrizations, we do not find these instances to be an impediment to mainstream use of gradient-based aerodynamic shape optimization of wings. The use of proper physics modeling and the application of practical design constraints enabled us to arrive at a single optimal design for the problem considered in this chapter. The insights gleaned from these results will help other researchers gain a better understanding of the wing design problem and formulate their problems to avoid possible sources of multimodality. Our experience is that in practical design problems, the number of local minima is automatically reduced merely by the standard constraints and requirements of the design problem. As stated in the introduction, we did not attempt to, nor did we, conduct an exhaustive search for all of the local minima in the design space. Our study confirms our opinion that it is not necessary or economical to engage in such activity for each new design problem.

The problems analyzed in this chapter only consider aerodynamic effects. Real-world wing design lies at the intersection of multiple disciplines, including aerodynamics, structures, and controls, among others. High-fidelity aerostructural analysis captures the interactions between aerodynamics and structures, the two disciplines with the strongest coupling in wing design. The use of high-fidelity aerostructural analysis in the optimization problem will introduce new design tradeoffs and have different effects on multimodality. Chapter 4 is an investigation of problems considering aerostructural effects.

Chapter 4

Exploration of Aerostructural Wing Design

In this chapter we investigate some aspects of transonic wing design through aerostructural optimization. Since transonic flow conditions differ significantly from the subsonic flow considered in Chapter 3, we start by re-evaluating the effect of sweep in the optimization problem. Then we demonstrate the capability of the MACH framework to explore the aerostructural design space by optimizing a rectangular wing in order to arrive at a swept wing typical of the transonic regime. We report the benefits of using a multi-level optimization acceleration technique for aerostructural optimization. Finally, we analyze the robustness of an optimized wing and propose a new method for preserving low-speed aerodynamic performance while optimizing for a cruise design condition.

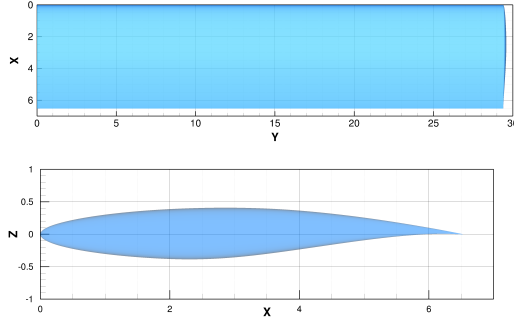
4.1 Problem Description

The CRM was designed as a benchmark for the verification and validation of CFD solvers across industry and academia⁹³. Subsequently, the CRM wing was adopted as the test case for a series of benchmark aerodynamic shape optimization problems⁵⁸. More recently, the undeflected (jig) shape of the CRM was reverse engineered as a benchmark for aerostructural analysis and optimization. In this work, we start with a rectangular wing and solve the same optimization problem defined for the undeflected Common Research Model (uCRM) in Brooks et al.⁵⁴. We created a rectangular wing with the same reference area and aspect ratio as the uCRM-9. The cross-section of the wing is the RAE-2822 airfoil with a trailing edge thickness of 5 mm. The planform and cross-section of the wing are depicted in Figure 4.1 and the initial geometric properties are given in Table 4.1.

The objective of the optimization problem is to minimize fuel burn for a specified

Table 4.1: Rectangular wing specifications

Property	Value	Units
Reference area	383.12	m ²
Half-span	29.38	m
Aspect ratio	9.01	
Mean aerodynamic chord	6.52	m
Sweep	0	degrees

**Figure 4.1:** Rectangular wing definition**Table 4.2:** Aircraft specifications

Property	Description	Value	Units
W_{fixed}	Fixed mass	100000	kg
W_{payload}	Payload	34000	kg
W_{reserve}	Reserve fuel	15000	kg
R	Mission range	7250	nm
c_t	TSFC	0.53	hr ⁻¹

mission length. The optimizer is free to change the shape of the wing and the sizing of the wingbox in order to achieve minimum fuel burn. Fuel burn is calculated in the following manner:

$$\begin{aligned}
 W_3 &= W_{\text{fixed}} + W_{\text{wing}} + W_{\text{payload}} + W_{\text{reserve}} \\
 W_2 &= W_3 \exp\left(\frac{R c_t C_D}{V C_L}\right) \\
 W_{\text{fuel}} &= W_2 - W_3
 \end{aligned} \tag{4.1}$$

The parameters W_{wing} , C_L , and C_D are subject to change during the optimization; fixed properties are defined in Table 4.2 and V is based on the Mach number and altitude.

Table 4.3: Grid dimensions

Label	N_{edge}	N_{chord}	N_{span}	$N_{\text{off-wall}}$	N_{total}	Max y^+
B2	2	22	33	32	52,096	2.15
A2	3	32	48	44	152,064	1.42
B1	4	44	66	64	416,768	1.19
A1	6	64	96	88	1,216,512	1.12
B0	8	88	132	128	3,334,144	1.15
A0	12	128	192	176	9,732,096	1.15

4.2 Methods

4.2.1 Computational Framework

We use the MACH framework to solve the optimization problem introduced above. For this study, we consider the solutions of both the multidisciplinary analysis and the coupled adjoint sufficiently converged when the L2 norm of the residual has decreased by 1×10^{-5} .

4.2.2 Preprocessing

CFD Meshing

We created two series of three meshes, for a total of six meshes, shown in Figure 4.2. The finest mesh (A0) was created first and then coarsened by a factor of $\sqrt{2}$ to get the starting mesh for the B-series (B0). The level 1 and 2 surface meshes in each series were coarsened by factors of 2 and 4, respectively, from the level 0 mesh. All surface meshes were extruded to a far-field distance of $100 \times \text{MAC} = 652 \text{ m}$ with an initial off-wall spacing of $5.74 \times 10^{-6} \text{ m}$. For these meshes, x is the streamwise direction and z is the lift direction. The dimensions of the grids are given in Table 4.3.

FEA Meshing

We patterned the structural mesh after the uCRM-9 wingbox. The fore and aft spars are located at 10% and 65% chord, respectively. The wingbox has 46 ribs, extending from the symmetry plane to the wingtip. Figure 4.3 shows the placement of the wingbox within the wing and the discretization of the wingbox. The structural mesh consists of 16,672 2nd-order quadrilateral shell elements.

Each component of the wingbox model is divided into 7 panels in the chordwise sense and 45 panels in the spanwise sense. Each of these panels consists of multiple shell elements. The panels of the skins and spars are modeled using smeared stiffeners, as described in Kennedy and Martins⁹⁴. The panels can also be grouped into design variable groups so that they share the same values for these parameters.




Figure 4.2: CFD meshes of baseline wing.




Figure 4.3: The structural mesh has 16,672 2nd-order shell elements.

Geometric parametrization

The geometry is parametrized using an FFD volume, as described in Section 2.2.1. We use a coarse FFD for the optimizations on the B2 and A2 grids and a fine FFD on the finer grids. The coarse FFD has 5 control points distributed with cosine spacing along the chord and 9 control points distributed evenly along the span, making a total of 90 FFD points (including top and bottom). The fine FFD has double the number of chordwise control points, for a total of 180. The coarse FFD is needed on the coarse CFD grids to maintain an appropriate ratio of CFD points to FFD points. As this ratio decreases, the optimizer has more and more control over the individual CFD points and is likely to produce a non-smooth surface. Generally we try to maintain a ratio of 4 CFD points to a single FFD point. The distribution of the control points in relation to the wing can be seen in Figure 4.4. The global taper and sweep




Figure 4.4: Geometric parametrization for rectangular wing optimization problem. On the left, the wing surface embedded in the FFD volume. On the right, the CRM planform is reproduced by modifying the FFD control points.

Table 4.4: Flow conditions for high-fidelity analyses

Flight condition	Mach	Altitude (ft)
Nominal cruise	0.85	37,000
2.5g pull-up maneuver	0.64	0
1.3g cruise buffet	0.85	37,000

design variables can be modified to convert the rectangular wing to a planform resembling the CRM wing.

4.3 General Optimization Problem

This section contains the description of a general optimization problem which is the basis for all of the results in this chapter. In the following sections, any deviations from or additions to the problem will be expressly stated. The optimization problem is defined in Table 4.5. Bounds and scaling factors for the variables and constraints are listed in the columns to the left, where applicable. The optimization problem requires three high-fidelity analyses: 1) nominal cruise, 2) 2.5g pull-up maneuver, and 3) 1.3g cruise buffet. The flow conditions for these three cases are listed in Table 4.4.

4.3.1 Objective function

The objective function is calculated with Equation 4.1. As mentioned previously, the optimizer can change W_{wing} , C_L , and C_D in order to decrease W_{fuel} . The wing mass is made up of the mass of the finite element model and an additional component to account for any increase in wing area.

$$W_{\text{wing}} = 2.5 W_{\text{wingbox}} + 4000 \frac{S_{\text{ref}}}{S_{\text{ref,orig}}} \quad (4.2)$$

Table 4.5: Rectangular wing aerostructural optimization problem description.

		Quantity	Lower	Upper	Scaling	
minimize with respect to	W_{fuel}	1				
	Angle-of-attack	3	0°	10°	0.1	
	Twist	7	-10°	10°	0.05	
	Sweep	1	0 m	25 m	0.01	
	Chord scaling	3	0.25	2.0	0.1	
	Sectional shape	180	-50 cm	50 cm	1	
	Panel thickness	131	2 mm	20 cm	100	
	Stiffener thickness	108	2 mm	20 cm	100	
	Stiffener height	91	5 mm	10 cm	100	
	Stiffener pitch	3	10 cm	30 cm	100	
	Panel length	108				
	Fuel tractions	301				
	Total fuel mass	1				
	Total number of design variables	937				
subject to	<i>Nonlinear constraints</i>	$L_{\text{nominal}} = (W_2 + W_3)/2$	1	0	0	1e-6
		$L_{2.5g} = 2.5W_2$	1	0	0	1e-6
		$C_{L,\text{buffet}} = 1.3(C_{L,\text{nominal}} + 0.05)$	1	0	0	1e-6
		Structural failure constraints	5		1	1
		Buffet-onset constraint	1		0.04	100
		$S_{\text{ref}} - S_{\text{ref,orig}}$	1	0	0	0.1
		Minimum wingtip thickness	15	10%		1
		Minimum trailing edge thickness	15	100%		1
		Minimum spar height thickness	30	60%		1
		Total fuel mass constraint	1			
		Fuel volume constraint	1			
		Fuel traction consistency constraints	301			
		Panel length consistency constraints	108			
	<i>Linear constraints</i>	LE/TE constraints	18			
		Monotonic constraint on chord scaling	2			
		$t_{\text{stiff},i} - t_{\text{panel},i}$	108	-2 mm	2 mm	
		$h_{\text{stiff},i} - t_{\text{stiff},i}$	108	0		
		$t_{\text{panel},i} - t_{\text{panel},i+1}$	104	-2.5 mm	2.5 mm	
$t_{\text{stiff},i} - t_{\text{stiff},i+1}$	104	-2.5 mm	2.5 mm			
$h_{\text{stiff},i} - h_{\text{stiff},i+1}$	88	-5 mm	5 mm			
	Total number of design constraints	1013				

In order to account for the lack of the fuselage and other surfaces, we make a few modifications to the lift and drag coefficients calculated by ADflow.

$$C_{L,\text{total}} = \frac{2.2 L_{\text{wing}}}{S_{\text{ref}} q} \quad (4.3)$$

$$C_{D,\text{total}} = \frac{2 D_{\text{wing}} + 3.453q}{S_{\text{ref}} q} \quad (4.4)$$

L_{wing} and D_{wing} refer to the lift and drag generated by the half-body CFD model, thus the factor of 2 in Equations 4.3 and 4.4. The drag area of 3.453 m² in Equation 4.4 is the product of the baseline reference area and the sum of the drag coefficients for the fuselage, empennage, and nacelle surfaces. The nacelle and vertical stabilizer combined contribute 30 drag counts, the same value used in the uCRM optimizations. To calculate the drag markup for the fuselage and horizontal stabilizer, we first

compute the ratio of the grid-converged drag values from the wing-only CRM case⁵⁸ and the DPW 4 wing-body-tail geometry. This ratio is then used to calculate the fraction of the uCRM drag pertaining to the fuselage and horizontal stabilizer. The complete calculation is as follows:

$$C_{D,FH} = \left(1 - \frac{C_{D,W,CRM}}{C_{D,WFH,DPW4}}\right) C_{D,WFH,uCRM} \approx 0.006 \quad (4.5)$$

where W, F, and H refer to wing, fuselage, and horizontal stabilizer, respectively. These modifications are used on all of the high-fidelity analyses.

4.3.2 Design variables

Each case has an angle-of-attack variable to enable matching the lift constraint. There are seven twist variables, each controlling the rotation of one of the spanwise FFD sections about the leading edge. The first two sections are fixed at zero twist. The sweep variable corresponds to the streamwise displacement of the wingtip leading edge. All other FFD sections are displaced linearly to create a straight leading edge. Chord scaling is controlled at FFD sections 1 (symmetry), 4, and 9. The intervening spanwise sections are scaled linearly to ensure straight leading and trailing edges. An example of the sweep and chord scaling variables is shown in Figure 4.4. As explained in Section 4.2.2, we use an FFD with 90 control points for the coarse CFD meshes and an FFD with 180 control points for the fine meshes. At each spanwise section, the FFD control points are restricted to in-plane displacements that are perpendicular to the freestream. The FFD control points regulate the cross-sectional shape of the wing.

The structure is divided into 131 design variable groups: 23 for the ribs, 18 for the spars, and 45 groups each for the upper and lower skins. Each of the rib design variable groups has a single panel thickness variable because the ribs are not modeled with the smeared stiffeners. The spar design variable groups share a single variable for stiffener height and another for stiffener pitch. Each spar group has its own variables for panel thickness and stiffener thickness. A single stiffener pitch variable is shared by all upper skin groups, and another is shared by all lower skin groups. All skin design variable groups have their own variables for panel thickness, stiffener thickness, and stiffener height. Finally, each of the design variable groups with smeared stiffeners (skins and spars) has a variable for panel length. The alignment of the stiffeners on the panels is calculated based on the initial panel reference axis, but does not change in the course of the optimization. Thus, while the stiffness matrix is updated to reflect changes in the length of the stiffeners, it does not account for changes in the orientation of the stiffeners. For optimization problems that allow the wing sweep to change, the stiffeners will remain aligned with the initial sweep. This creates an artificial benefit for the wing sweep to remain

close to the initial value, because the structure is most efficient when the stiffeners are aligned with the wing sweep. For small variations in wing sweep, the effect of this discrepancy is minor, however, for the rectangular wing case, we expect large changes in wing sweep. We have two ways of managing this issue. First, for optimization problems that start from the baseline rectangular wing, we manually set the stiffener orientation to a value that is close to the expected optimal wing sweep. Second, we use a multi-level approach to optimization, in which successive optimizations start from where the previous one left off (see Section 4.5).

The weight of the fuel in the wing is applied as a uniform traction to the lower skin of the wingbox between the symmetry plane and the 44th rib. This region of the lower skin is made up of 301 panels, each of which receives a design variable for the weight of fuel to be applied as a traction. Additionally there is a variable for the total amount of fuel in each case. The cruise and buffet cases have a fuel load corresponding to the mid-cruise point while the maneuver case has a fuel load corresponding to the full fuel load.

4.3.3 Design constraints

Each of the analysis points has its own lift constraint. The structural failure constraints are computed by aggregating a failure criterion over a group of elements using the KS aggregation method. There are two types of failure criterion added as constraints. The stress failure criterion is calculated as the von Mises stress in an element multiplied by a safety factor and divided by the material yield strength. The buckling failure criterion is based on the critical buckling load of a stiffened quadrilateral panel. We include three stress failure constraints: one for the ribs and spars, one for the upper skin, and one for the lower skin. Two buckling failure constraints are added: one for the ribs and spars, and one for the upper skin. The buffet-onset constraint was developed and tested by Kenway and Martins⁹⁵. In summary, buffet onset is correlated to the percentage of the upper skin with separated flow. They determined that the threshold of separated flow should be 4% of the area of the upper skin. This constraint is added to the 1.3g cruise case, so as to maintain the required 30% margin to buffet during cruise.

A constraint is added to preserve the original reference area. Thickness constraints are added along the wingtip to prevent excessive flattening of the wingtip cap, which could crush cells in the volume mesh. These constraints limit the thickness of the wingtip cap to a minimum of 10% the original thickness. Any decrease of the trailing edge thickness is prevented with a set of 15 thickness constraints at 99% chord. Thickness constraints are also added along the fore and aft spars to prevent decrease beyond 60% of the original value. This set of thickness constraints was added to prevent excessive thinning of the outboard wing (see Figure 4.11). A

set of constraints (dubbed LE/TE constraints) are added to the pairs of FFD control points at the leading and trailing edges of each section to ensure equal and opposite displacement. This ensures that the shape variables do not cause shearing twist, which would be redundant with the global twist variables. In Chapter 3, we motivate the use of a monotonic constraint on chord variables to ensure that the chord decreases monotonically from the root to the wingtip. In this work, we include a linear constraint on the chord scaling variables to enforce this property.

There are two fuel load constraints to ensure that the total fuel load variable for each case is consistent with the actual amount of fuel being carried by the aircraft. Each of the design variable groups with smeared stiffeners has a linear constraint to maintain a difference of less than 2 mm between the panel thickness variable and the stiffener thickness variable. There are also linear adjacency constraints to limit the difference in stiffener height, stiffener thickness, and panel thickness between adjacent panels to 1 cm, 5 mm, and 5 mm, respectively. There are 64 nonlinear constraints added to ensure the deformed panel lengths are consistent with the 64 panel length variables. An additional 602 nonlinear constraints exist to ensure that the fuel traction variables are set to the correct value. Finally, for each fuel load, there is a constraint to ensure that the total fuel volume can fit inside the wingbox.

4.3.4 Preliminary Structural Optimization

Initially, all of the structural members have uniform thickness and stiffener sizings. We can start the aerostructural optimization from a more reasonably sized structure if we first optimize the structural sizing with a set of fixed aerodynamic loads. This structural pre-optimization is a cycle with five iterations. In each iteration, we run an aerostructural analysis with the current structural sizing in order to get the aerodynamic loads. Then we apply the aerodynamic loads to the structure and run an optimization to minimize the structural mass with respect to the failure constraints on the 2.5g maneuver. We repeat this process five times in order to arrive at a semi-converged aerostructural state. For the plank geometry, the structural pre-optimization produces a wingbox weight (W_{wingbox}) of 11,874 kg, which corresponds to a total wing weight (W_{wing}) of 37,686 kg. The optimized structure for the CRM-shaped planform has a wingbox weight of 10,968 kg and a total wing weight of 35,954 kg.

4.4 Aerodynamic Shape Optimization

Before addressing the aerostructural optimization problem presented in Table 4.5, we take a moment to analyze a few aspects of the transonic ASO problem. In Section 3.4, we showed that for subsonic flow, the design is relatively insensitive to

Table 4.6: Aerodynamic shape optimization problem description.

		Quantity	Lower	Upper	Scaling
minimize with respect to	C_D	1			
	Angle-of-attack	1	0	10	0.1
	Twist	7	-10	10	0.05
	Chord scaling	3	0.25	2.0	0.1
	Sweep	1	0 m	25 m	0.01
	Sectional shape	90	-0.5	0.5	1
	Total number of design variables	102			
subject to	$C_L = 0.5$	1	0	0	1e-6
	$S_{\text{ref}} - S_{\text{ref,orig}}$	1	0	0	0.1
	Minimum wing thickness	400	25%		1
	Minimum wingtip thickness	15	10%		1
	Minimum trailing edge thickness	15	100%		1
	Minimum t/c along spars	30	60%		1
	LE/TE constraints	18			
	Monotonic constraint on chord scaling	2			
	Minimum volume constraint	1	100%		
		Total number of design constraints	958		

the sweep distribution. In transonic flow conditions, wave drag is a major factor and wing sweep is a desirable feature. Wing sweep has the effect of decreasing t/c from the perspective of the flow normal to the leading edge of the wing, which in turn reduces wave drag. Wave drag can be reduced by sweeping the wing backward or forward, making the optimization multimodal. We demonstrate this by optimizing a wing at various fixed sweep angles. The optimization problem description for this study is found in Table 4.6. These optimizations used the coarsest mesh (B2) because it was sufficient to reveal the desired trend. Figure 4.5 shows the variation of C_D , C_{M_y} , and α for the different optimized wings. As predicted, the variation of C_D is bell-shaped with the peak corresponding to a quarter-chord sweep of 0° . Interestingly, the optimal angle of attack decreases as the wing is swept forward.

Backward sweep tends to promote wingtip stall. With forward sweep, wingtip stall is delayed, preserving healthy flow across control surfaces. Despite the fact that forward sweep is more advantageous aerodynamically, there are additional aeroelastic effects that complicate the design of forward-swept wings. A conventional aft-swept wing exhibits negative bend-twist coupling which tends to unload the wingtip as it deflects under aerodynamic forces. A forward-swept wing has the opposite effect, leading to divergent behavior unless the wing is extremely stiff. Most aircraft have swept-back wings to reduce structural weight.

For the rest of the studies in this chapter, we restrict the tip sweep to positive values only in order to eliminate the multimodality observed in Figure 4.5. Before moving onto aerostructural optimizations, we investigate three more aerodynamic optimization problems. These optimization problems are based on Table 4.6, but differ in geometric parametrization and starting configuration. We use a multi-level optimization procedure, in which successive optimizations are run using increas-




Figure 4.5: Wave drag can be reduced by sweeping the wing forward or backward, resulting in a multimodal problem. Each of the plotted points is the result of an aerodynamic shape optimization at a fixed sweep angle.

ingly finer CFD grids. Each optimization starts from the final state of the previous optimization. This process saves computational time by converging most of the way to the optimum on coarser grids (see Section 4.5). The three problems are defined as follows:

- Case 1) Straight leading and trailing edges. Start from plank geometry.
- Case 2) Straight leading and trailing edges. Start from CRM planform.
- Case 3) Curved leading and trailing edges. Start from CRM planform.

The curved leading and trailing edges of case 3 are created with a combination of chord scaling and x -displacement. As explained in Section 2.2.1, these design variables are applied at the control points of an FFD reference axis. In this case, the reference axis is placed along the leading edge of the wing, with one control point at each of the spanwise FFD sections. Chord scaling is enabled at all nine reference axis control points and x -displacement is enabled at all but the root control point. These three optimization problems are solved using meshes B2 through A1.

Figure 4.6 shows the planforms of the three optimized wings. Case 1 does not sweep the wing backward because it gets stuck on the forward side of the curve in Figure 4.5. The FFD reference axis is located on the leading edge of the wing, so the chord scales about the leading edge. When the tip chord decreases in size

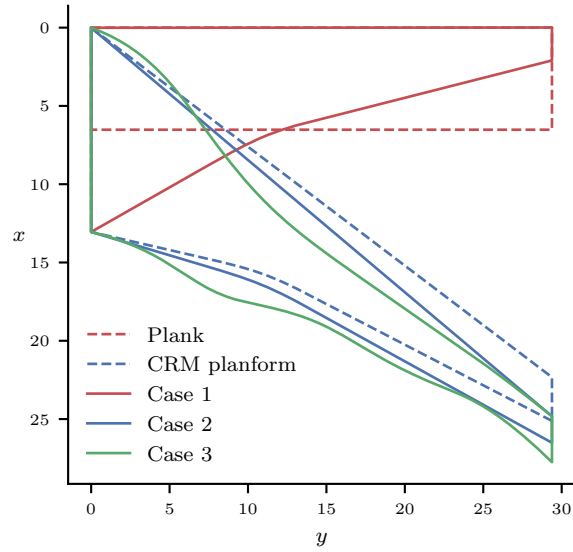


Figure 4.6: Planforms

the quarter-chord sweep becomes negative. Once the sweep becomes negative, the optimizer can see no benefit in sweeping it backward, so the sweep variable remains at the lower bound (which is the starting value for the plank geometry). The tip sweep in cases 2 and 3 goes to the upper limit of 25 meters, which is consistent with the downward slope of the C_D curve in Figure 4.5. It is possible that by relaxing the upper bound of the sweep variable, the optimizer would eventually locate the optimal sweep angle for the single-point case, but we were unable to test this hypothesis due to mesh warping restrictions. The spanwise properties and airfoil cross-sections of the optimized wings are provided in Figure 4.7. The clear differences in airfoil shape are indicative of the differences in planform shape. With the exception of the root and the tip sections, case 1 features prototypical supercritical airfoils—the optimizer’s best chance to compensate for the unswept planform. The bulging root section of case 1 is necessary to permit thin supercritical airfoils over the rest of the wing while still satisfying the fuel volume constraint. By contrast, the cross-sections of cases 2 and 3 exhibit a more conventional forward loading and camber distribution. The chord distributions of cases 1 and 2 are practically identical, whereas the chord distribution of case 3 is bell-shaped with a distinctive flare at the tip. The optimized drag counts for the three cases are 236.8, 224.3, and 223.3, in order of enumeration. Thus, the additional freedom present in the geometric parametrization of case 3 affords an additional drag reduction of approximately 1 drag count.




Figure 4.7: Comparison of three aerodynamic shape optimizations.

4.5 Multi-level Optimization Procedure

In previous work on aerodynamic shape optimization by Lyu et al.⁵⁸, a sequence of optimizations were performed on progressively finer meshes in order to reduce the total computational time of the optimization. Each optimization starts where the previous optimization ends, so that the maximum benefit is extracted from each mesh. The coarser meshes do not resolve the flow as accurately, but they still can provide sensitivities that point the optimizer in the direction of the true optimum. With each successive mesh level, the optimizer finds itself closer to the optimum with an enriched set of sensitivities to guide its path. For this study, we adopt the multi-level optimization method with the three coarsest meshes: B2, A2, and B1. The purpose of this study is to compare the multi-level approach to an optimization using only the B1 mesh. The optimization problem is a simplified version of the one laid out in Table 4.5, where the sweep is fixed at 34.9° . Thickness constraints along the spars are not included in this comparison.

The result of this comparison is shown in Figure 4.8. Both methods converge to basically the same shape, but the computational cost of the multi-level approach is 60% of the cost of the single-level optimization. The direct comparison does not tell the whole story here though. In practice, the process of setting up and troubleshooting a new optimization problem generally requires countless debugging runs. This author has found great value in having a very coarse mesh handy during

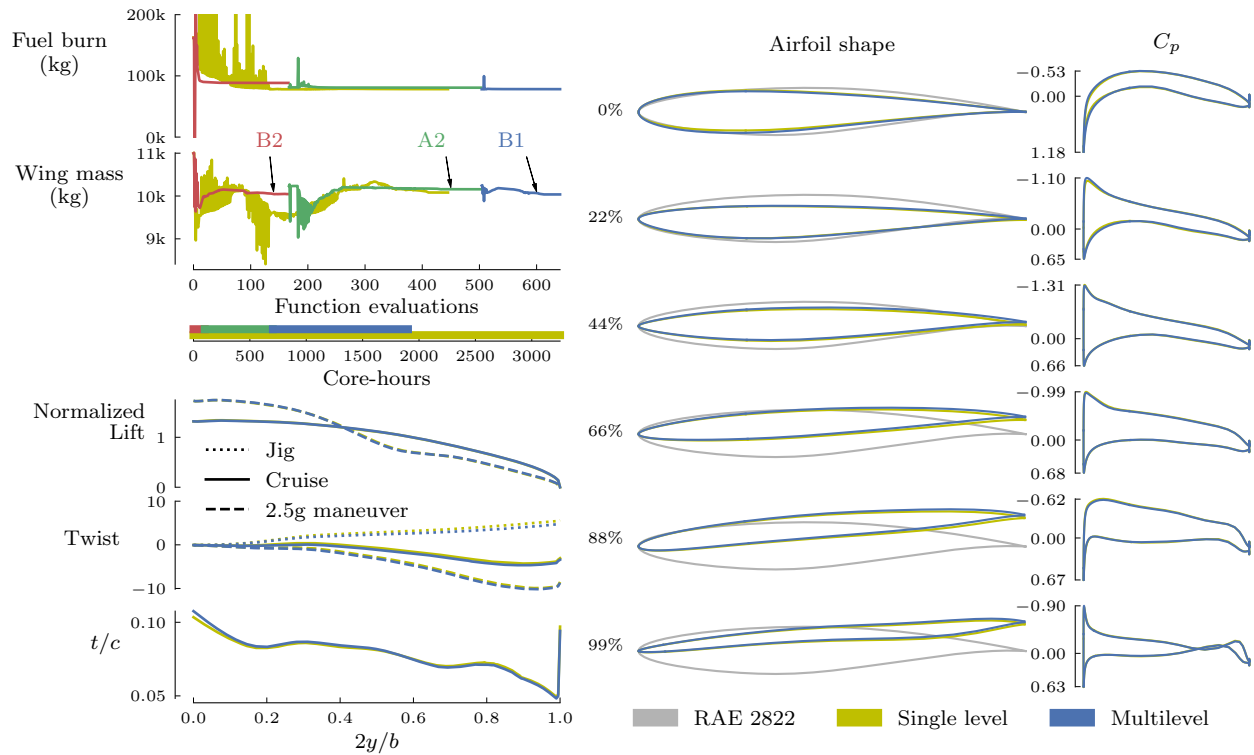


Figure 4.8: For this problem, the multi-level approach achieves basically the same design as the single-level optimization on the finest grid—at 60% of the computational cost.

this initial phase. The cost savings due to troubleshooting with a coarse mesh are not easily accounted for, but probably far outweigh the cost savings due to a single optimization run.

4.6 Single-point Optimization

Now we look at the full optimization problem described in Table 4.5. For this problem, we use grid levels B2, A2, B1, and A1 successively to arrive at the final result. The coarse FFD (90 control points) is used on grid levels B2 and A2 and the fine FFD (180 control points) is used on grid levels B1 and A1. The optimized wings from each grid level are shown in Figure 4.9. The planform of the starting geometry is shown in grey.

The transonic flow condition presents an interesting tradeoff between wave drag and wing mass in relation to the sweep variable. To decrease wave drag, the optimizer can increase wing sweep or modify airfoil shape. However, increasing the wing sweep necessitates a heavier wing to support increased bending loads. On the B2 wing, the tip sweep reaches the upper limit. The airfoils of the A2 wing are very similar to the B2 wing, and yet the sweep decreases by 2° , indicating that the improved resolution of the flow field favors less sweep. The final two optimizations

Table 4.7: Results of single-point optimization.

Case	Grid level	FFD	W_{fuel} (kg)	W_{wing} (kg)	L/D	Sweep (deg)
Starting from plank	B2	Coarse	88,740	37,531	20.3	37.3
	A2	Coarse	80,727	35,758	21.7	35.2
	B1	Fine	77,243	32,725	22.3	36.0
	A1	Fine	75,834	31,919	22.5	36.7
Starting from CRM	A1	Fine	75,667	31,002	22.5	36.7
Coarse FFD	A1	Coarse	76,941	33,027	22.4	33.1
No shape variables	A1	Fine	90,928	39,167	20.0	38.1




Figure 4.9: Each successive level of mesh refinement yields additional design changes.

converge to a sweep value in between the first two. Switching to the fine FFD allows the optimizer to fine-tune the airfoil shape for the single design point, resulting in a fairly sharp leading edge and a pronounced suction peak. A distinctive lower surface concavity forms at the leading edge which is reminiscent of the result of an airfoil optimization of the RAE-2822 by Drela⁶³. The optimizations on the finer grids are also able to produce more passive load alleviation, as shown in the difference between the load and twist distribution at the cruise and maneuver design points.

Figure 4.10 shows that the improvements achieved from each successive optimization are preserved in a grid refinement study. The optimized wings from each of the four grid levels are analyzed with the finer grids. The flow condition for this comparison is the cruise condition at $C_L = 0.5$. We chose $C_L = 0.5$ because each of the wings was optimized to a lift coefficient near 0.5. The baseline grid convergence




Figure 4.10: Drag convergence study for the baseline and single-point optimized wings. For the baseline wing, the drag increases as the mesh is refined because the shock is resolved more accurately on finer grids. The improvements on the optimized wings are preserved as the mesh is refined.

is also plotted to show the increase in drag as the large shock structure is captured more accurately on successively finer meshes. This comparison gives an indication of the value added for each successive optimization in the multi-level approach. In comparison to the B2 optimum, meaningful gains are realized in the A2 and B1 optimizations. However, the drag reduction going from the B1 optimum to the A1 optimum is marginal. The minimal differences in the wing design between the B1 and A1 optima suggest that there is no need to continue the multi-level optimization onto the next grid level.

In addition to starting the optimization with the plank geometry, we also ran the optimization using the CRM planform as the initial design. In addition to the differences in starting planform, the initial structural sizing for each of these runs differs due to the structural pre-optimization (see Section 4.3.4). As shown in Figure 4.11, the optimizer converges to nearly the same design starting from the plank geometry and starting from the CRM planform. A closer look at the differences in final design variables reveals that while the optimizer did converge to the same general shape, there are significant differences between the two optimized designs, especially with regard to structural sizing (Figure 4.12). Notably, the design variables of larger magnitude are more likely to converge to the same value, whereas smaller design variables exhibit greater variance.

Figure 4.11 also shows the effect of the number of shape variables on the overall design. When the coarse FFD is used, the wing has reduced sweep and is unable to produce as much passive load alleviation as the wing optimized with the fine FFD.




Figure 4.11: Comparison of single-point optimized designs. The optimizations starting from the plank and CRM planforms converge to nearly the same design.

Additionally, the optimizer is unable to tailor the leading edge radius as precisely with the coarser FFD. The wing optimized without shape variables converges to a planform design with a constant chord on the inboard section of the wing. Normally, the optimizer can reduce wave drag by increasing chord while keeping the thickness constant, but in this case, the optimizer avoids increasing the root chord because it has no control over the thickness ratio of the wing. As shown in Table 4.7, including shape variables in the optimization reduces the objective by 14,000 kg. Further fuel burn reduction of over 1,000 kg is realized by using the fine FFD over the coarse FFD. This improvement is likely due more to the 2,000 kg reduction in wing weight than the marginal improvement in L/D .

4.6.1 Synopsis

In this section, we investigated several aspects of the single-point aerostructural optimization problem. We compared results obtained using various CFD grid levels and saw that the final cross-sectional shape, planform, and structural sizing are all affected by the grid level. This comparison also revealed that there are diminishing returns on using increasingly finer CFD grids. We showed that the optimizer was able to converge to a similar design from two radically different starting points, allaying concerns that gradient-based optimization is not suitable for design space

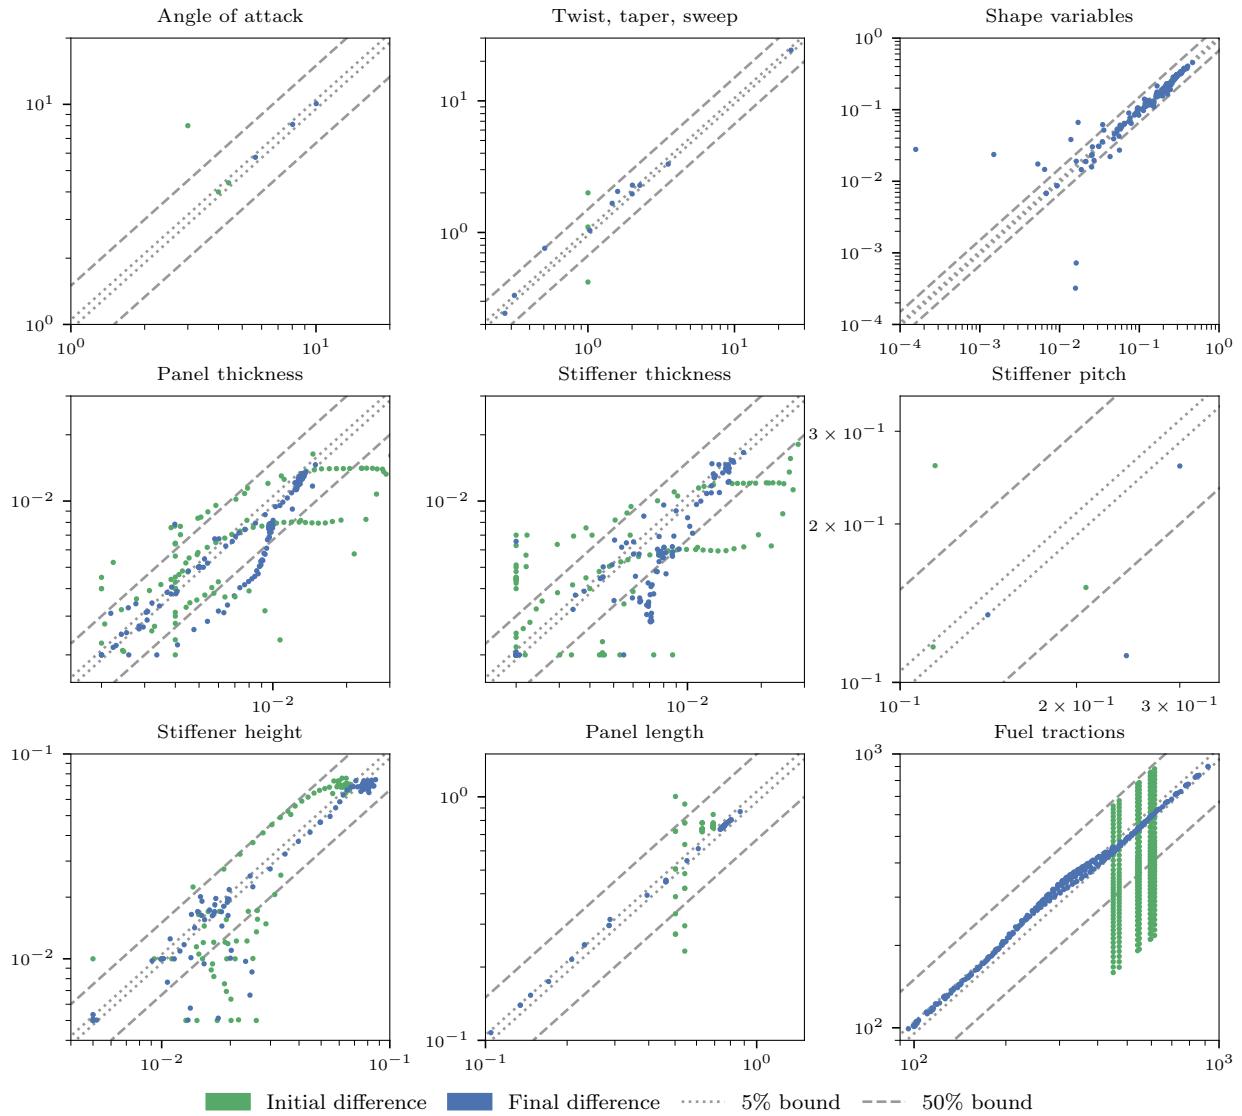


Figure 4.12: The difference in initial and final design variables for wing optimization problems starting from a plank platform and a CRM platform.

exploration. Finally, we considered the impact of the number of shape variables on the overall design and found that the airfoil shape is implicitly linked to the planform shape and the structural sizing.

4.7 Robust Design Optimization

4.7.1 Background

Single-point optimizations are prone to exhibit poor off-design performance. One of the most common solutions for this problem is to set an objective function that is a weighted average of the performance at multiple design points. The set of design conditions included in the objective is referred to as a *multipoint* stencil. Thus, even though the optimization problems solved in Section 4.6 included cruise, maneuver, and buffet analysis points, we will designate them as single-point designs because the objective function was only based on a single design point. Using a multipoint objective improves the average performance across the stencil at the expense of the nominal design point. However, it can result in intermittent performance, wherein the design functions optimally at the specified design conditions but poorly in the intervals. Drela⁶³ reported this phenomenon in a set of airfoil optimization studies and showed that increasing the number of points in the stencil helped to curb this tendency. In wing ASO, Lyu et al.⁵⁸ obtained a more robust design using a 5-point stencil than with a single-point optimization. The multipoint design had a weak shock across the stencil, whereas the single-point design had completely eliminated the shock at the nominal design point. They also found that the multipoint design had a larger leading edge radius than the single-point design. Kenway and Martins⁶⁴ compared different multipoint stencils of varying size and composition and found a good compromise between robustness and computational expense with a carefully chosen 5-point stencil. Multipoint optimization has also been demonstrated in the context of aerostructural wing design^{52,54,96}.

However, while there has been extensive comparison of single-point and multipoint designs with ASO, the same cannot be said for aerostructural wing optimization. Additionally, most of the past efforts on multipoint design have focused on robust cruise performance without considering the impact of design changes at low-speed, high-lift conditions. Preserving low-speed, high-lift performance in a wing optimization problem is notoriously difficult due to the complications that arise from modeling and parametrizing high-lift devices. The difficulties associated with high-lift devices can be avoided by considering clean wing performance at low-speed, high-lift conditions. To this end, Wakayama and Kroo⁹⁷ and Ning and Kroo³⁸ have shown that constraining $C_{L,max}$ using critical section theory results in a more practical planform design. In airfoil optimization, Buckley et al.⁹⁸

added a constraint on $C_{l,\max}$ into the multipoint objective function in order to meet safety requirements at a low-speed condition. Rather than constraining $C_{L,\max}$, Khosravi and Zingg⁵⁵ simply included climb drag in the multipoint objective function to encourage improvement in that regime. The whole issue is often skirted by simply imposing limitations on the geometric parametrization to prevent changes that would adversely affect high-lift performance (e.g. minimum leading edge thickness).

4.7.2 Setting the stage

In this section we evaluate the impact of using a multipoint objective in aerostuctural wing design optimization. Sophisticated methods of determining a representative set of design conditions and weights have been researched by Toal and Keane⁹⁹ and Liem et al.⁹⁶. However, in this work we use a simple 3-point stencil in C_L -space based on Case 4.2 from Kenway and Martins⁶⁴. The nominal C_L is calculated based on the mid-cruise weight at Mach 0.85 and 37,000 ft, just as in the single-point optimization problem. The auxiliary design points are analyzed at $C_{L,\text{nominal}} \pm 0.05$ at the same Mach number and cruise altitude. Initially, we tried $C_{L,\text{nominal}} \pm 0.025$, following Brooks et al.⁵⁴, but found little variation between the single-point and multipoint designs. The three design points are weighted equally, so the objective function is simply the arithmetic mean of the fuel burn calculated from L/D at each flow condition. In all other respects, the multipoint optimization is identical to the problem described in Table 4.5.

We also experiment with a novel approach for ensuring airworthiness at low-speed, high-lift conditions. Using publicly available flight data of the Boeing 777-200ER (an aircraft with similar specifications to the CRM), we determined Mach 0.4 and 10,000 ft to be a low-speed flight condition that should exhibit good aerodynamic performance. This flight condition is shown in relation to the flight data in Figure 4.13. As will be shown, both the single-point and multipoint designs exhibit severe separated flow when evaluated at this flight condition at a high angle of attack. By contrast, a single-point optimization with fixed RAE-2822 cross-sections performs well. Instead of minimizing the drag at this flight condition, we propose using a constraint to limit separation on the wing. Our approach is inspired by buffet-onset constraint of Kenway and Martins⁹⁵ and uses the same formulation for the separation sensor. Based on the analysis of the single and multipoint wings, the constraint allows no more than 10% of the upper surface of the wing to have separated flow. The analysis point for the separation evaluation is constrained to generate enough lift to sustain the nominal takeoff weight (W_2) of the aircraft at Mach 0.4 and 10,000 ft. Both the chosen flight condition and the value for the upper limit of the separation constraint are somewhat arbitrary and are subject to




Figure 4.13: Analysis point for low-speed, high-lift separation constraint is placed at the boundary of the climb profile for the Boeing 777-200ER.

change for other applications. This work is mainly concerned with introducing the application of a low-speed separation constraint and evaluating its impact on the optimized design. The correlation of this constraint with established airworthiness regulations remains for future work.

The main objective of this study is to evaluate the impact of two different robust optimization approaches: 1) a multipoint objective and 2) a separation constraint at a low-speed, high-lift flight condition. Here is a summary of the designs considered in this section:

- (SP) The single-point design from Section 4.6 (started from the CRM planform).
- (MP) A three-point design with $C_{L,nominal}$, $C_{L,nominal}-0.05$, and $C_{L,nominal}+0.05$ at Mach 0.85 and 37,000 ft.
- (SP-LS) A single-point design with a separation constraint at Mach 0.4 and 10,000 ft.
- (SP-NSV) Result of single-point optimization without shape variables (RAE-2822 cross-sections).

4.7.3 Results

While only the MP optimization problem included the multipoint stencil in the objective function, the other three optimized designs were analyzed at the multipoint flight conditions for the purposes of comparison. Characteristics of the four optimized designs are given in Table 4.8 The multipoint design achieves the best




Figure 4.14: Optimizing for robust performance in cruise *and* climb incurs ~1000 kg increase in cruise fuel burn.




Figure 4.15: The multipoint design is very similar to the single-point design, but including the low-speed separation constraint elicits striking modifications.

average fuel burn, but does slightly worse than the single-point design at the nominal flight condition. The SP-LS design burns more fuel across the three points than the SP (+1.0%) and MP (+1.6%) designs, but compared to the SP-NSV design it is 16.0% more fuel efficient. The relative performance of the four designs is shown in Figure 4.14.

Surprisingly, the airfoils of the multipoint design are nearly identical to the single-point design (Figure 4.15). Increased sweep on the multipoint design is the major geometric difference between the two. Accordingly, the multipoint wing is

slightly heavier than the single-point wing, with most of the weight gain in the skins and the aft spar. By contrast, the SP-LS wing is substantially different than the single-point or multipoint wings. One of the most striking differences is the disappearance of the distinctive concavity at the leading edge of the lower surface. The SP-LS cross-sections have a larger leading edge radius as a result and seem more typical of a traditional airfoil. The C_p curves feature reduced suction peaks (due to larger leading edge curvature) and more aft loading than the other designs. Sweep is reduced by 5° compared with the single-point design. The distribution of t/c is significantly lower from 40-80% span, likely decreasing wave drag, which would otherwise increase due to the reduction in sweep. Passive load alleviation at the 2.5g maneuver condition is severely degraded for the SP-LS wing. The variation in spanwise loading between the cruise and maneuver conditions is minimal, as opposed to the SP and MP designs where the cruise loading is elliptical and the maneuver loading is bell-shaped. This means that the SP-LS wing requires a heavier structure in order to achieve an elliptical cruise lift distribution and still satisfy the failure constraints at the maneuver condition. The upper and lower skins see the greatest increase in weight, but the fore spar is also significantly heavier than the single-point design.

Given the differences in weight and geometry, the cruise drag polars for the SP, MP, and SP-LS designs are surprisingly similar (Figure 4.16). These three designs have nearly the same performance at the nominal cruise point; the variation in nominal fuel burn is due to the differences in weight rather than the aerodynamic efficiency. Moving outward from the nominal design point in either direction, the multipoint design is the most robust, followed by the SP-LS design. The multipoint design does especially well as C_L is increased from the nominal design point. Most of the improvement in the average fuel burn over the single-point design comes from the $C_{L,nominal} + 0.05$ design point. The SP-LS design achieves higher L/D than the SP design at all three design points yet burns more fuel due to a heavier structure.

The right side of Figure 4.16 shows the performance of the four designs at the climb flight condition. The four designs were analyzed across a sweep of angle of attack at increments of 1° . The SP and MP designs do not converge after the angle of attack exceeded 9° due to massive separation on the upper surface. Separated flow leads to a sharp increase in C_D as angle of attack is increased for both of these designs. The trend of the drag polars indicates that the $C_{L,max}$ for these wings would be lower than the C_L required to satisfy the lift constraint $L = W_2$. The multipoint design fares slightly better than the single-point design, but both fail to meet expected airworthiness at the climb condition. The drag polar of the SP-LS wing mimics that of the SP-NSV wing up until 9° angle of attack. At that point

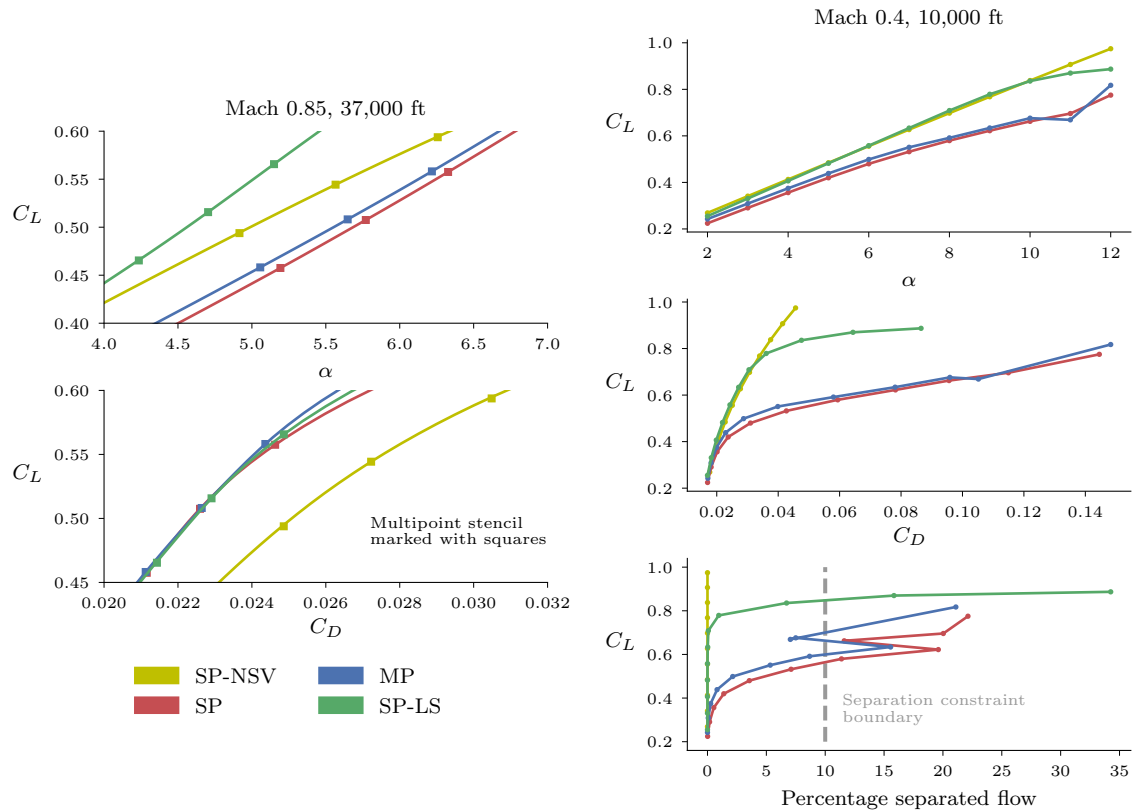


Figure 4.16: The low-speed separation constraint improves robustness at both cruise and climb flight conditions.

separation ensues, but the 10% threshold is not exceeded until nearly 11° angle of attack. At 9° angle of attack, the SP-LS wing generates 24% more lift and 54% less drag than the SP wing. The near-complete elimination of separated flow on the SP-LS wing at 9° angle of attack is impressive (Figure 4.17).

It is worth noting that the buffet-onset constraint is inactive at the end of the SP-LS optimization. This suggests that low-speed separation is a more stringent requirement than buffet-onset. Conversely, the buffet-onset constraint was active for both the single-point and multipoint designs, but did not improve the low-speed, high-lift performance of these wings. One possible reason for this finding is that the low-speed separation constraint has the effect of decreasing the cruise angle of attack. Separation at the 1.3g flight condition is more likely to occur if the wing is cruising at a higher angle of attack, as is the case for the SP and MP designs.

4.7.4 Synopsis

Multipoint objective functions typically consider design points in the cruise regime. While this does result in more robust cruise performance, we have shown that cruise-optimized designs do not, in general, perform well at off-design (i.e. climb)

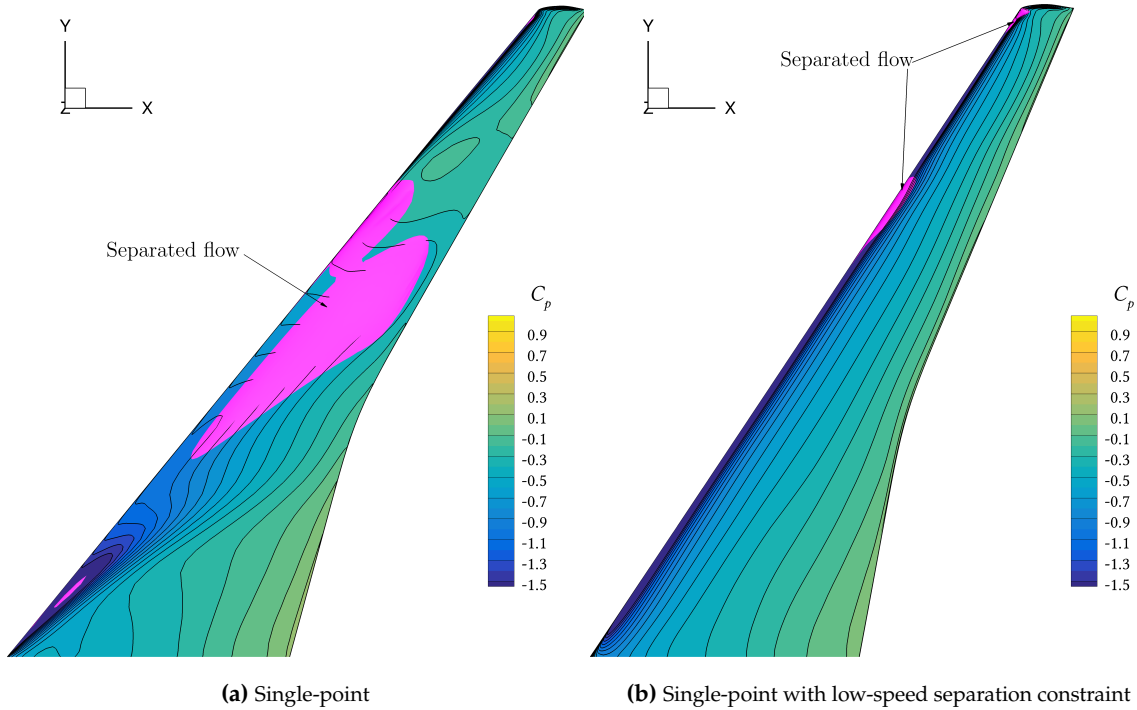


Figure 4.17: Separation on the upper surface at 9 degrees angle of attack is nearly eliminated with the low-speed separation constraint.

conditions. We have shown that the inclusion of a low-speed separation constraint in the SP-LS case pushes the optimizer to a completely different design than the single-point and multipoint problems. The single-point, multipoint, and SP-LS designs do not differ significantly in terms of aerodynamic performance at cruise. Rather, the changes in the wing design seem to have the greatest impact on performance at off-design conditions. The low-speed separation constraint dramatically improves the high-lift capability of the wing at low-speed conditions. Additionally, the design changes required to satisfy the separation constraint have the effect of reducing passive load alleviation at the 2.5g maneuver loading, resulting in a substantial increase in structural weight. The differences in fuel burn among the designs are more related to varying structural weights than significant stratification of cruise performance. Importantly, the single-point and multiple designs appear to be unviable concepts due to their poor low-speed, high-lift characteristics, whereas the SP-LS inspires much more confidence as a practical design.

4.8 Summary

In this chapter we show that a gradient-based optimizer can be used to explore the aerostructural wing design space. We experiment with a new method for improving the robustness of an optimized design at low-speed, high-lift conditions. This

Table 4.8: Optimization results.

Property	Units	SP-NSV	SP	MP	SP-LS
Combined fuel burn	kg	92,145	76,592	76,261	77,518
W_3	kg	188,167	180,001	180,191	183,002
Wing weight	kg	39,167	31,001	31,191	34,001
Upper skin	kg	5,040	3,957	3,976	4,529
Lower skin	kg	4,902	3,643	3,637	4,238
Ribs	kg	1,691	1,234	1,260	1,235
Fore spar	kg	409	165	187	217
Aft spar	kg	423	199	214	180
Sweep	deg	38.1	36.7	37.4	30.8
<i>Nominal</i>					
Fuel burn	kg	90,928	75,667	75,817	77,518
Angle of attack	deg	5.57	5.77	5.65	4.70
C_L		0.5443	0.5075	0.5082	0.5157
C_D	cts	272.3	226.0	226.4	229.1
L/D		19.99	22.46	22.44	22.51
$C_{L,nominal} - 0.05$					
Fuel burn	kg	91,642	79,147	78,972	80,074
Angle of attack	deg	4.92	5.19	5.06	4.24
C_L		0.4938	0.4575	0.4582	0.4654
C_D	cts	248.6	211.6	211.3	214.3
L/D		19.86	21.63	21.69	21.72
$C_{L,nominal} + 0.05$					
Fuel burn	kg	93,866	74,963	73,995	75,765
Angle of attack	deg	6.26	6.33	6.22	5.15
C_L		0.5937	0.5575	0.5581	0.5657
C_D	cts	304.9	246.3	243.6	248.6
L/D		19.48	22.64	22.91	22.75

method involves setting a constraint to limit flow separation at a flow condition representative of the climb phase. The constraint significantly affects the optimized design, indicating that post-optimization modifications to a cruise-optimized design would be unlikely to produce an optimal design. Without the constraint, both single-point and multipoint optimized designs are not airworthy at the climb condition. This underscores the importance of factoring off-design robustness into the optimization problem when designing wings.

Chapter 5

Practical Wing Design Optimization

Traditionally, the planform, airfoil shape, and structure of a wing are designed separately. One of the first steps is to determine the area and basic shape of the wing planform that will support the mission requirements. Historical data and low-fidelity models can be used to obtain this basic wing definition rapidly. Eventually, the in-flight wing shape is designed by aerodynamicists to achieve the desired performance. It is then the job of structural engineers to design a jig shape and structure that will revert to the flying shape when deformed aeroelastically in flight. The wing design is then transferred back and forth between the aerodynamics and structures groups to converge on the lightest wing that can satisfy performance requirements. While this description is over-simplified, it highlights the separation that exists between the design of the planform, wing shape, and structure. The complexity inherent to wing design has led to this segmented approach. Advances in computing power and numerical methods have enabled MDAO techniques. One of the aims of MDAO is to remove artificial barriers in the design process that may have arisen in the past due to organizational practicalities. When the entire design can be robustly optimized simultaneously, it should be possible to converge to a better design with a significant reduction in time and resources.

MDO of wings has been an active area of research for many years, especially with low-fidelity aerostructural models^{92,100,101}. However, only recently has high-fidelity aerostructural wing optimization reached a level of maturity at which realistic wing design can be done. Some examples include optimization of the CRM⁵⁴ and the D8¹⁰².

Despite the important work in the academic literature, competition in industry makes it difficult to determine to what extent these methods have been adopted and borne fruit. In an effort to bridge this gap between industry and academia, we

collaborated with Embraer to optimize the wing design of an authentic regional aircraft using the MACH framework. This collaboration allows us to evaluate the benefit that could be obtained by applying MDO in the industrial aircraft design process. In the results presented herein, we compare the baseline wing to a wing with optimized structural sizing, airfoil shape, and twist distribution. This comparison serves as a litmus test for the optimization problem—significant deviation from the baseline design would indicate a failure to properly capture the design requirements. Then we extend the optimization problem to include planform variables, allowing the optimizer more freedom to trade between structural weight and aerodynamic performance. For each of these wing optimization problems we model the aerodynamics of the full aircraft geometry, including fuselage, wing, and empennage. In this way, subtle interactions between the wing and the rest of the geometry are not neglected, and we can ensure that the optimized design is properly trimmed. Additionally, we are able to experiment with sizing the horizontal stabilizer to account for changes in the wing planform.

In the spirit of practical wing design, it is important to choose an objective function that suitably quantifies the actual performance of the aircraft. Fuel burn is often used as an objective function in wing design optimization problems because it represents operating cost. Various methods exist for estimating fuel burn, each with assumptions and approximations that influence the outcome of the optimization. For example, the fuel burn calculation can be based on a single flow condition, or it can aggregate information from multiple design points. As explained in Section 4.7, using numerical optimization to design a wing for a single flight condition can have undesired consequences at off-design conditions⁶³. A multipoint optimization problem, in which the performance at multiple flight conditions is incorporated into the objective function, can produce a more robust design.

There are different ways to formulate a multipoint objective function. Commonly, the objective function is a weighted average of a figure of merit computed at each flight condition. However, the weighting has a significant effect on the end result and it is not always clear how to choose the weights a priori. Liem et al.⁹⁶ devised a method to calculate the weights for 25 different flight conditions in order to achieve minimum fuel burn for the CRM geometry. They generated a surrogate model of aerodynamic coefficients with aerostructural analyses of the baseline geometry and then calculated the fuel burn for hundreds of representative missions using the surrogate model. Finally, they computed the sensitivity of the total fuel burn with respect to each of the 25 flight conditions and used this as the weight in the objective function for the aerostructural optimization. Not surprisingly, the aircraft optimized with this multi-mission objective function weighting performed much better than a single-point optimum when tested on the full spectrum of

Table 5.1: R-jet aircraft Specifications

Description	Symbol	Units	Value
Wingspan	b	m	28.72
Reference area	S	m ²	95.4
Aspect ratio	\mathcal{R}		8.6
Maximum takeoff weight	MTOW	kg	51,800
Operating empty weight	OEW	kg	27,900

missions.

In the study just cited, all of the flight conditions included in the fuel burn aggregation were sampled from the cruise flight phase. For an aircraft that typically flies long-range missions, the fuel burn in climb and descent is negligible. However, a regional aircraft regularly flies missions in which the fuel used in climb and descent exceeds that used in cruise. Therefore, if fuel burn is the quantitative measure of aircraft performance, it is important to include design points from the climb and descent segments in the fuel burn aggregation. Liem et al.⁸⁰ addressed this concern in a paper similarly focused on regional aircraft design optimization. They optimized the mission profiles for two missions of varying range by integrating fuel burn across the climb, cruise, and descent segments. Then they used the optimized mission profiles to compute fuel burn while optimizing the aerostructural design of the aircraft wing. In the present work, we also account for the fuel burn in climb and descent, albeit in a slightly different way. Furthermore, we take a closer look at how the design is impacted by the consideration of full mission fuel burn as opposed to only considering cruise.

5.1 Problem Description

5.1.1 Geometric Modeling

Basic specifications of the Embraer regional jet are provided in Table 5.1. Embraer furnished the CAD descriptions of the fuselage, wing, and empennage surfaces that define the outer mold line (OML). Each component of the geometry is meshed individually using structured hexahedral cells and then combined into a single overset mesh (Figure 5.1). The grid used for optimization has just over 2 million cells. The finite element model of the wingbox is based on a structural layout from Embraer. The wingbox mesh consists of 8,018 quadrilateral cells which are modeled as second-order mixed interpolation of tensorial components (MITC) shell elements. Together, the OML and wingbox form the baseline design of the aircraft, which, for the purposes of this study, has been designated R-jet.

We also created a supplementary wing-only model of the aircraft called R-wing.




Figure 5.1: Computational grids for the R-jet configuration.

Table 5.2: The R-wing mesh yields a 1500x decrease in proc-hours for a single aerostructural analysis when compared with the R-jet overset mesh.

Geometry	Cells	N Cores	Solution time	Adjoint time
R-jet	2,071,850	79	1261 s	587 s
R-wing	30,528	1	87 s	77 s

This wing geometry has the same general specifications detailed in Table 5.1, but it differs from the R-jet in that it has an RAE-2822 cross-section, a straight trailing edge, and a conventional wingtip (as opposed to a winglet). The R-wing wingbox has the same basic layout as the R-jet, but the ribs are all aligned with the freestream. Table 5.2 shows the size of the R-wing CFD mesh and the substantial reduction in computational expense when compared with the full R-jet overset mesh. The R-wing mesh reduces the proc-hours by 1500 and 1000 times for analysis and a single adjoint, respectively. The primary purpose for creating the R-wing geometry was to enable rapid optimization cycles for studying the effects of mission analysis on the optimized design. However, a side benefit of having an inexpensive computational model is that it accelerates the inevitable troubleshooting process for a brand new optimization problem. Generally, optimizer parameters and scaling factors tuned on the simple model should be applicable to the more expensive model.




Figure 5.2: Computational grids for the R-wing configuration juxtaposed with the R-jet grids.

5.1.2 Mission Description

The objective of the design problem is to minimize the average fuel burn of three representative missions, weighted as follows:

$$\overline{W}_{\text{fuel}} = 0.5W_{\text{fuel,NOM}} + 0.2W_{\text{fuel,LR}} + 0.3W_{\text{fuel,HS}} \quad (5.1)$$

Details of the three missions can be found in Table 5.3. The total fuel burn for each of these missions is the difference in weight from takeoff (W_0) to landing (W_5), where the basic mission profile is defined as indicated in Figure 5.3. It is expected that most of the fuel is burned in the climb and cruise segments, and that the ratio of climb fuel burn to cruise fuel burn will decrease as the mission range increases. For instance, in Liem et al.⁸⁰ the climb and cruise segments accounted for 32% and 26% of the total fuel burn for a 500 nm mission and 13% and 68% of the total for a 2000 nm mission. For aircraft that typically fly missions where the distance covered in climb is small compared to the total mission range, the benefit of incorporating the climb segment into the total fuel burn objective is negligible. In such cases, it would still be important to consider flightworthiness requirements of the climb segment as constraints in the optimization problem (as shown, for example, in Section 4.7). However, as the ratio of climb range to cruise range increases, it becomes more important to consider the fuel burned in climb, especially because the fuel burn per mile is greater in climb than in cruise. It is well known that more fuel is burned in climb than in cruise for an equal distance covered. To simplify

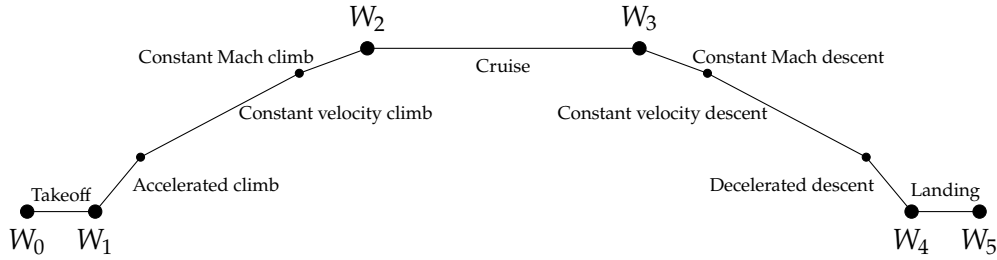


Figure 5.3: Basic mission profile.

Table 5.3: Mission details

Code	Description	Range (nm)	Mach	Altitude (ft)	Payload (kg)
NOM	Nominal	600	0.78	36,000	10,000
LR	Long-range	1000	0.74	36,000	10,000
HS	High-speed	300	0.80	36,000	10,000

the estimation of climb fuel burn, the hypothetical cruise fuel burn over the climb distance is augmented with a climb fuel increment factor, which is computed from empirical data sets^{103,104}.

It is instructive to visualize the variation in the relationship between the climb and cruise segments with respect to mission range, payload, and cruise altitude. Figure 5.4 was created by computing the fuel burn over the climb and cruise segments for missions ranging from 300 to 2000 nm with varying payloads and cruise altitudes. The data was generated using the mission analysis, engine model, and aerodynamic model described in Section 5.2, applied to the R-jet aircraft. The first two rows of Figure 5.4 show the ratios of climb fuel to cruise fuel and climb range to cruise range, respectively. As expected, both ratios increase exponentially as the total mission range decreases. Also, the ratios increase with increasing payload and cruise altitude. At any given data point, the magnitude of the fuel ratio is more than twice that of the range ratio, confirming that the fuel burn per mile is greater in climb than in cruise. The inverse of the rate of fuel burn per distance covered, known as specific range, is plotted in the final row of Figure 5.4. Here again, we see that the specific range in cruise is more than twice that achieved in climb. In all of these plots, the data corresponding to the R-jet payload and cruise altitude are highlighted in blue. From this preliminary study, it is apparent that for the 600 nm nominal mission and especially the 300 nm high-speed mission, it is critical to include the climb fuel burn in the objective function. In Section 5.3, we show how the optimized design is impacted by this modeling decision.

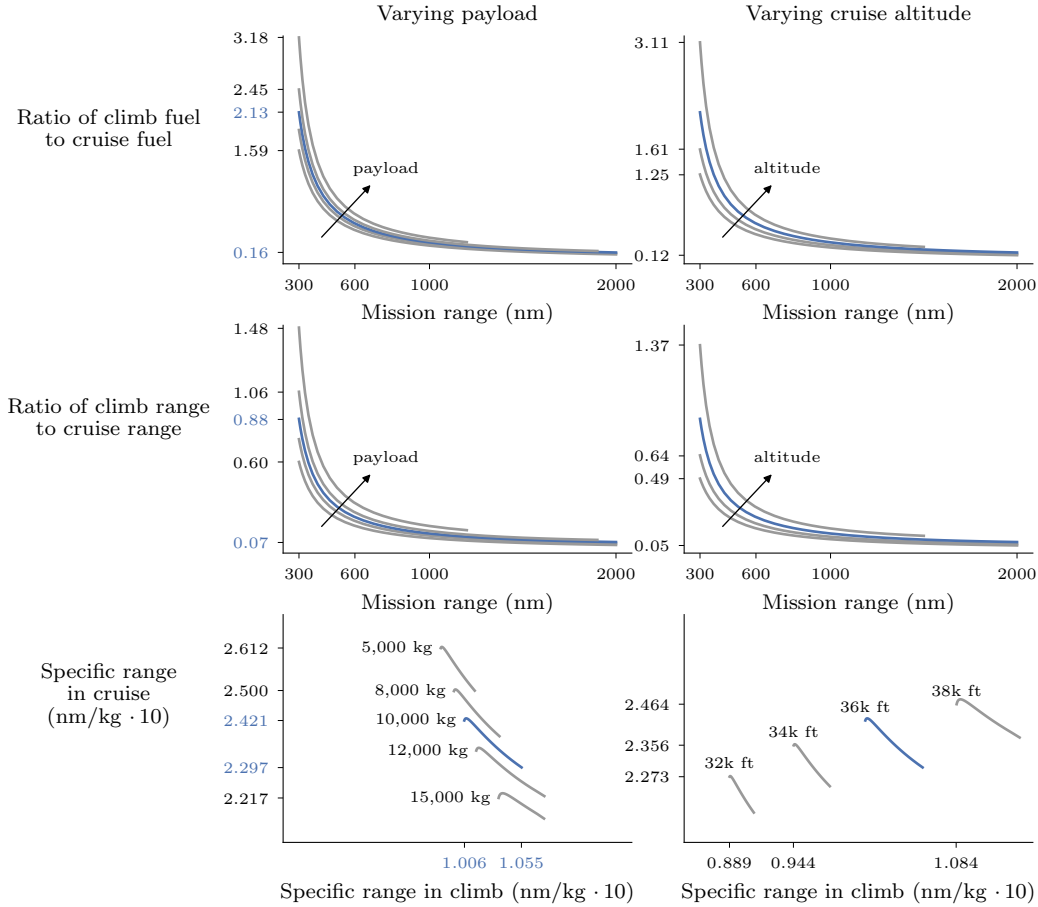


Figure 5.4: The ratio of climb fuel burn to cruise fuel burn grows exponentially as mission range decreases. The left column shows data for five different payload weights and the right column shows data for four different cruise altitudes. The data corresponding to the R-jet aircraft description is highlighted in blue.

5.2 Computational Framework

The results in this chapter are obtained using the MACH framework, which is described in Chapter 2. In this section, we elaborate on additional details of the computational framework that are specific to the R-jet problem.

5.2.1 Mission Analysis

For the mission profile depicted in Figure 5.3 the formula to compute fuel burn is

$$W_{\text{fuel}} = W_5 \left(\frac{W_0}{W_1} \frac{W_1}{W_2} \frac{W_2}{W_3} \frac{W_3}{W_4} \frac{W_4}{W_5} - 1 \right) \quad (5.2)$$

The final weight of the aircraft, W_5 , is the sum of the operating empty weight (OEW), payload, and reserve fuel. It varies during the optimization due to changes

Table 5.4: Four methods for calculating the fuel burn and range of a mission.

	Full-Bréguet	Full Mission	Cruise-Bréguet	Hybrid
W_0/W_1	0.98	0.98	0.98	0.98
W_1/W_2	1.0	MissionAnalysis	0.977/0.976/0.977	MissionAnalysis
W_2/W_3	Bréguet	MissionAnalysis	Bréguet	Bréguet
W_3/W_4	1.0	MissionAnalysis	0.993/0.993/0.992	MissionAnalysis
W_4/W_5	0.995	0.995	0.995	0.995
Climb range (nm)	0	MissionAnalysis	110/116/105	MissionAnalysis
Cruise range (nm)	600/1000/300	MissionAnalysis	373/769/67	MissionAnalysis
Descent range (nm)	0	MissionAnalysis	117/115/128	MissionAnalysis

in the structural sizing and wing shape. The fuel fractions in Equation 5.2 can be calculated in various ways. At the most basic level, they can be taken from historical data. In this work, the takeoff (W_0/W_1) and landing (W_4/W_5) fuel fractions are fixed at 0.98 and 0.995, respectively. The climb and descent fuel burn can also be represented with fixed values. In Chapter 4, the climb and descent fuel fractions are implicitly set to unity, while Equation 2.4 is used to approximate fuel burn over the entire mission range. This approximation is appropriate for medium to long-haul flights (Figure 5.4), and is often used in the literature. Alternatively, the climb and descent fuel fractions can be set to historical values, and the cruise fuel burn can be computed based on a restricted range corresponding to the estimated distance covered in cruise flight.

For a more general solution, the fuel burn and distance of the climb, cruise, and descent segments can be integrated using MissionAnalysis. As described in Section 2.4, this entails splitting the mission profile into multiple intervals and analyzing the aerodynamic and engine performance at the interval endpoints. For example, the data for Figure 5.4 was generated using 31 intervals for climb, 15 intervals for cruise, and 18 intervals for descent. The climb and descent profiles are split into subregions characterized by accelerated (decelerated), constant calibrated airspeed (CAS), and constant Mach flight (Figure 5.3). The integration intervals are apportioned to the various subregions to achieve an acceptable level of resolution across the mission profile. The mission is solved simultaneously by driving the weight differences between endpoints of consecutive intervals to zero. Thousands of aerodynamic analyses are required to converge to a valid solution, and thousands more are necessary to calculate sensitivities. The cost of using CFD for these aerodynamic analyses would be exorbitant. Liem et al.⁸⁰ used CFD to generate a surrogate model of aerodynamic properties at each optimization iteration. The surrogate model was then used to efficiently process the aerodynamic analyses required to solve the mission problem. In this work, we use a conceptual-level aerodynamic model provided by pyConcept to reduce the computational cost of the mission analysis.

As stated in the chapter introduction, one of the aims of this study is to evaluate the impact of using different fuel burn computation methods on the optimized wing design. The four methods considered are listed in Table 5.4. Whenever three values are given, separated by slashes, they correspond to the nominal, long-range, and high-speed missions of Table 5.3. For all methods, MissionAnalysis uses aerodynamic data from pyConcept whereas the Bréguet equation is always used in conjunction with high-fidelity aerostructural solutions from MACH. Thus, the Hybrid method uses a combination of low-fidelity data for climb and descent and high-fidelity data for cruise. The fixed fuel fractions and ranges for the Cruise-Bréguet method are the values calculated for the baseline configuration using the Full Mission method. It is worth noting that none of these methods in and of themselves are novel contributions. Some of them are frequently used in the literature and have already been demonstrated in this dissertation. For example, the Full-Bréguet method was used in Chapter 4 and the Full Mission method was used to analyze the missions for Figure 5.4. The unique contribution of this study lies in comparing the methods to each other, with a special focus on how they impact the final design in a wing optimization problem.

5.2.2 Aerodynamic Analysis

As mentioned previously, both high-fidelity and low-fidelity aerodynamics models are used in the optimization problems presented in this chapter. The high-fidelity data are produced using the coupled aerostructural analysis described in Section 2.3. To account for the lack of fuselage and empennage in the R-wing configuration, we apply an additive drag markup of $D/q = 0.75 \text{ m}^2$ and a multiplicative lift markup of 1.05. With these markups included, the mission fuel burn computed for the R-wing configuration is similar to that of the R-jet configuration. The computational grids for the aerodynamics and structures are shown in Figure 5.1 for R-jet and Figure 5.2 for R-wing. By contrast, the conceptual-level formulas in pyConcept require only a few data points representing the basic shape of the wing geometry. The required points are extracted from slices of the R-wing and R-jet wing surfaces. For a given cross-section slice, pyConcept uses points on the leading and trailing edges and a point on each of the upper and lower surfaces at the location of maximum thickness (Figure 5.5). With these points, pyConcept can compute necessary quantities such as planform area, quarter-chord sweep, and the value and location of the MAC. As explained in Section 2.4.1, these points are embedded in the FFD so that they are updated as the design changes during the optimization.

On the baseline configuration, the drag estimation from pyConcept matches well with the high-fidelity aerostructural result at the cruise flow condition (Figure 5.6). As the speed and altitude of the flow condition decrease, the low-fidelity analysis

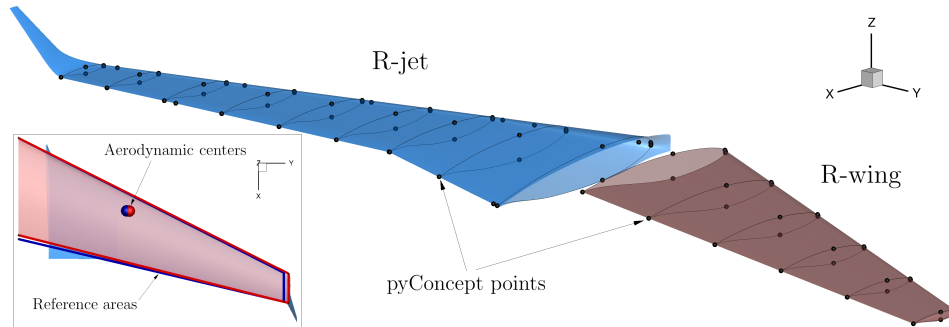


Figure 5.5: Points are extracted from the R-wing and R-jet wings in order to create a basic representation of the geometry in pyConcept.

underestimates drag by as much as 20%. Importantly, the difference between low-fidelity and high-fidelity analyses on an optimized design is less than or equal to that of the baseline design. We note that the terms high-fidelity and low-fidelity in this instance are used in a relative sense. The CFD mesh used for the R-wing analysis is relatively coarse and is only used to obtain rough estimates of aerodynamic performance. However, it does provide richer information than the pyConcept model.

5.2.3 Engine Model

The engine performance is represented using a surrogate model built from engine data of thrust and fuel consumption with respect to true airspeed, altitude, and temperature. We use a surrogate model from the Surrogate Modeling Toolbox (SMT)* called Regularized minimal-energy tensor-product splines (RMTS)¹⁰⁵. SMT is especially suited to gradient-based optimization because it provides analytic gradients of the functions of interest with respect to the independent variables. The available thrust at a given flight condition is throttled to the value needed to satisfy minimum residual climb limits (in climb) or offset drag exactly (in cruise). The relationship between engine fuel consumption and throttle setting is modeled as a quartic polynomial based on engine data. All of the engine data were provided by Embraer.

5.3 Simple Wing Optimization

In this section, we compare three different methods of mission analysis and evaluate their impact on the optimization of the R-wing configuration. We first compare multipoint optimizations of the Cruise-Bréguet and Hybrid methods (considering

*<https://github.com/SMTorg/smt>

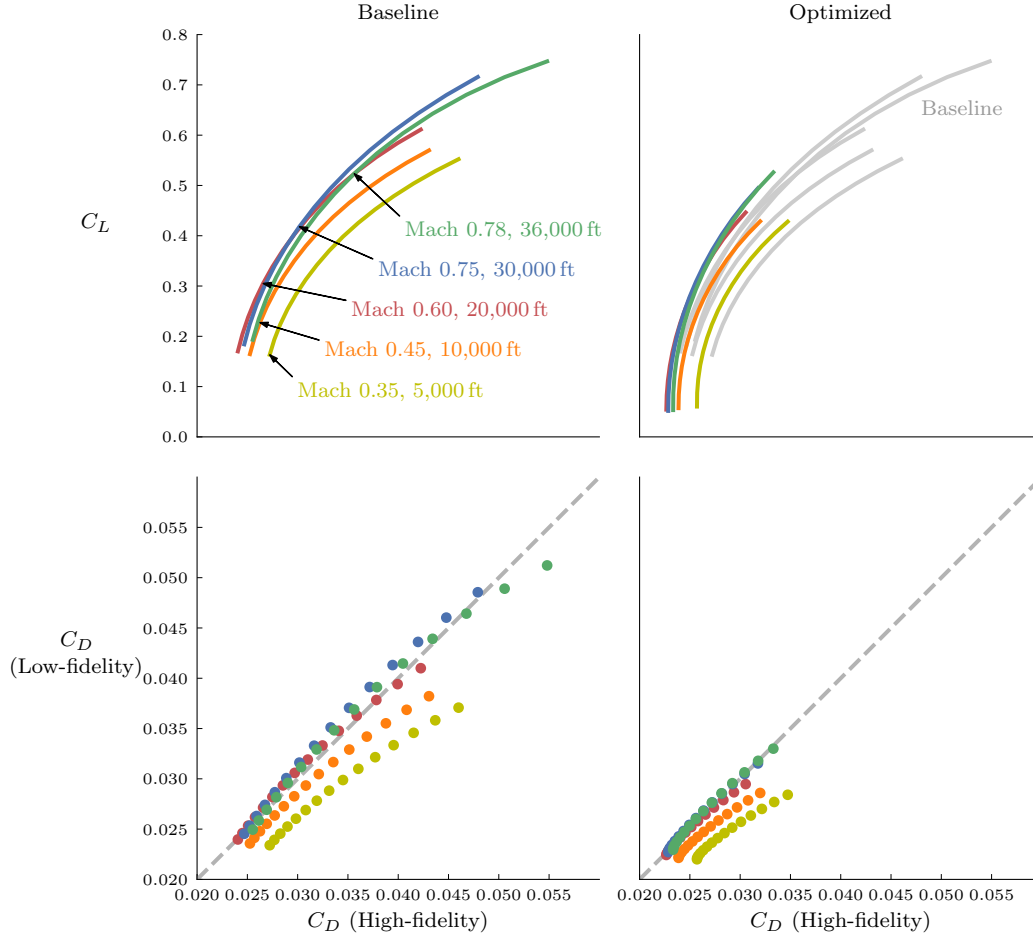


Figure 5.6: On both baseline and optimized designs, the low-fidelity drag polar from pyConcept coincides with the high-fidelity aerostructural result at cruise flow conditions. The error in the low-fidelity method grows as the speed and altitude decrease.

the three missions of Table 5.3). Then we conduct single-point optimizations of each mission individually using the former methods plus the Full-Bréguet method.

5.3.1 Optimization Problem

The R-wing optimization problem is given in Table 5.5. For this study, the aggregated fuel burn objective function defined by Equation 5.1 has been augmented with an additional term, operating empty weight, to create a multi-objective problem. The objective function is repeated here for clarity:

$$\text{minimize } \beta \overline{W}_{\text{fuel}} + (1 - \beta) \text{OEW} \quad (5.3)$$

Table 5.5: Description of the full R-wing optimization problem. The quantities for each design variable correspond to the green curve in Figure 5.7.

		Quantity
minimize	$\beta \overline{W}_{\text{fuel}} + (1 - \beta) \text{OEW}$	1
with respect to	Angle of attack	4
	Twist	4
	Planform variables (span, chord, sweep)	7
	Sectional shape	128
	Structural sizing	356
	Panel lengths	146
	Fuel load tractions	108
	Total fuel mass	4
	Mission fuel weights	12
subject to	Range constraints	3
	Trim constraints ($L = W$)	4
	Structural failure constraints	4
	Planform area cannot decrease	1
	Wing geometric constraints	76
	Miscellaneous mission constraints	159
	TACS linear constraints	420
	TACS nonlinear constraints	258
	Fuel load constraints	4

The parameter β is varied between 0.5 and 1.0 to generate a Pareto front that illustrates the tradeoff between aircraft weight and fuel burn. The Pareto front was not extended below $\beta = 0.5$ because it was apparent from the trend of the curve that further reductions in structural weight would incur an unacceptable increase in fuel burn. We also experiment with three different combinations of geometric design variables. This brings the final tally of optimization runs to 24: 2 mission analysis methods, 4 points along the Pareto front, and 3 different sets of design variables.

For each optimization problem there are four different high-fidelity analysis points, each of which has a design variable to control angle of attack. We consider three combinations of geometric variables

1. Only twist variables (at 4 spanwise stations)
2. Twist and planform variables (1 span, 1 sweep, and 5 chord scaling)
3. Twist, planform, and local shape variables

The structure is parametrized with the blade-stiffened panels described in Section 2.3.3. The number of structural design variables for each case remains the same. Fuel load traction variables are used to vary the tractions on the skin panels so that, in aggregate, they impart a load equal to the total fuel mass variable. The total fuel mass variable, in turn, is constrained to be consistent with the actual fuel burn of the given mission for the current iterate. For the optimizations using the Hybrid mission analysis, there are two variables corresponding to the initial and

final weights of the climb and descent segments of the three missions. For these studies we do not allow the cruise altitude to vary from 36,000 ft.

Each mission is constrained to achieve its specified range. The high-fidelity analysis point for each mission is constrained to generate enough lift to support the mid-cruise weight (the average of W_2 and W_3). Yield stress and buckling constraints are evaluated under the loads imposed by a 2.5g pull-up maneuver at Mach 0.734 and 15,000 ft. This design point also carries a payload of 13,000 kg and a full fuel load of 12,900 kg. Yield stress failure constraints are aggregated over the upper and lower skins separately. Buckling constraints are aggregated over the ribs and spars, and separately over the upper skin. The geometry is constrained so that the thickness cannot decrease at 1% and 99% chord. Additionally, at 60% chord the thickness cannot decrease below 80% of the originally value. The chord variables are constrained to decrease monotonically from root to tip. For the case with shape variables, constraints are applied to the FFD points at the leading and trailing edges to ensure that they move in equal and opposite directions to prevent shearing twist. We also place a constraint that the planform area cannot decrease when the planform variables are active. The Hybrid mission analysis method uses constraints to force the mission slack functions to be greater than or equal to zero (see Section 2.4). The linear constraints in TACS consist of adjacency constraints to prevent large differences in thickness between adjacent panels. The nonlinear constraints in TACS include compatibility constraints for the panel lengths and fuel tractions. Finally, as mentioned previously, each mission has a total fuel mass variable that must be constrained to match the fuel burn of the current iterate.

5.3.2 Results

Figure 5.7 shows the two sets of Pareto curves generated for the wing-only optimization. The design variables and constraints for each of the red, blue, and green curves are given in the final three columns of Table 5.5. Each Pareto front is created by varying β between 0.5 and 1. As expected, adding design variables to the problem shifts the Pareto front to the lower-left hand corner, signifying overall improvement in the design. Most of the benefit of additional geometric freedom comes in the form of reduced fuel burn, especially for higher values of β . Impressively, there is roughly a 1-2% reduction in fuel burn between each of the Pareto fronts.

Each point along the Pareto front is accompanied by a scaled outline of the planform and a number referring to the aspect ratio of the optimized wing. Interestingly, when only cruise fuel burn is considered in the optimization objective, the aspect ratios do not increase significantly beyond the baseline geometry, even when $\beta = 1$. However, when the full mission fuel burn (including climb and descent) is considered, the aspect ratios for higher β values increase significantly. The plots

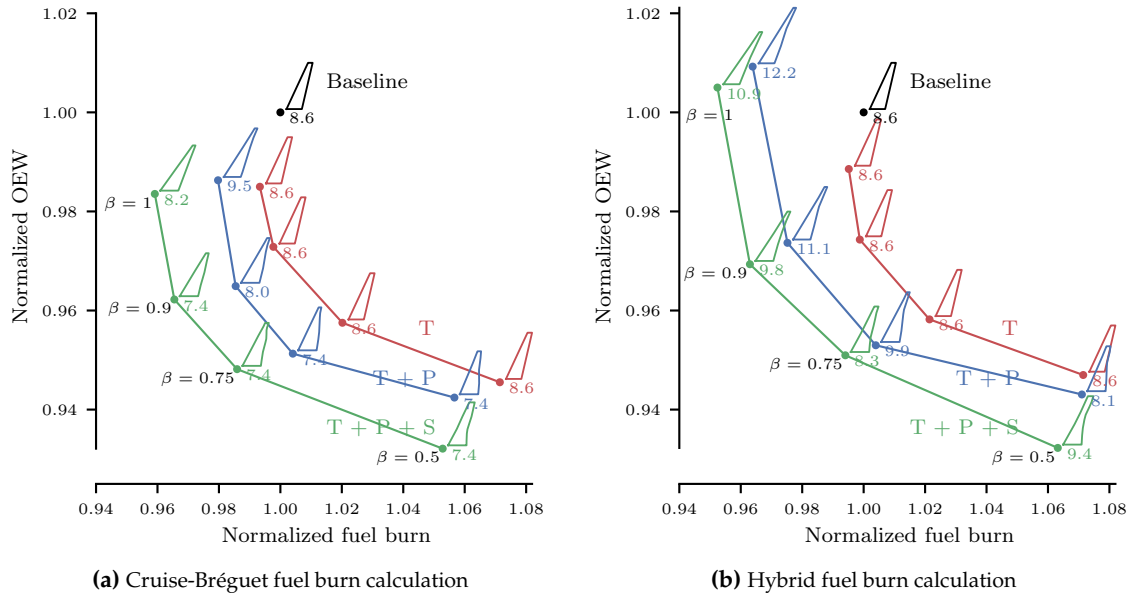


Figure 5.7: Pareto fronts from optimizations using different methods of fuel burn calculation.

in Figure 5.7 highlight these differences. The right-hand plot shows the expected trend, that the aspect ratio should increase in order to decrease fuel burn.

For a regional jet, the mass of fuel burned over a standard mission is on par with the mass of the wingbox structure. Therefore, the fuel burn reduction resulting from a decrease in wingbox mass is closer in magnitude to that achieved through aerodynamic enhancements than it would be for an aircraft with a longer mission. This seems to be the reason why accounting for the fuel burn over the entire mission causes an increase in aspect ratio for $\beta = 1$. When only the cruise segment is considered, the benefit of increasing aspect ratio does not outweigh the penalty in fuel burn due to the increase in structural mass required to support a higher aspect ratio. This result validates our use of mission analysis for the full configuration optimization study. Perhaps for a larger aircraft, constant fuel fractions would be sufficient, but for a regional jet, the gains made in fuel burn reduction over climb and descent really matter.

One possible explanation for the increased emphasis on drag reduction in the Hybrid method results is that the low-fidelity model used in the mission calculations is over-predicting the possible drag reduction on the climb and descent segments. On the contrary, post-optimization high-fidelity analysis of points along the climb and descent profiles yields greater drag reductions than were predicted by the low-fidelity analysis (Figure 5.8). Thus, it is more likely that using the low-fidelity aerodynamic analysis in the mission calculations yields a conservative result for the optimized design.




Figure 5.8: The low-fidelity aerodynamic analysis used in MissionAnalysis underestimates the drag reduction at several points along the climb and descent profiles.

5.3.3 Synopsis

In wing design, the optimal aspect ratio is dependent on the tradeoff between structural weight and fuel burn. The balance between structural weight and fuel burn can be adjusted explicitly by controlling their relative importance in a weighted average objective function. For instance, in our Pareto front studies, the wings optimized for minimum fuel burn ($\beta = 1$) had higher aspect ratios than those with more emphasis on structural weight ($\beta = 0.5$). On the other hand, the optimal aspect ratio is implicitly dependent on the relative magnitudes of the structural weight and fuel weight—specifically the quantities that are variable in the optimization. For a short-haul regional jet, the fuel and the wing structure have weights of similar magnitude. Thus, the assumptions built into the fuel burn calculation can have a significant impact on the balance between fuel weight and structural weight and, by extension, the aspect ratio of the optimized wing. Modeling the climb segment, and including its fuel burn in the objective function, increases the incentive for the optimizer to reduce fuel burn through drag reduction.

5.4 Full Aircraft Optimization

Now that we have demonstrated the value of including mission analysis in the optimization of a simple wing geometry, we will move on to optimizing the wing of the full R-jet configuration. The primary aim of this optimization problem is to see how the optimizer alters the design of the original regional jet. Since the baseline wing is already a “good” design, major differences between the baseline and optimized designs indicate either a shortcoming of the optimization problem or a previously unattainable improvement that is made possible with the MDO approach. We start by sizing the wingbox structure to achieve minimum mass while still satisfying the structural constraints for the baseline configuration. In this way, comparisons of structural weight, and as a result fuel weight, between the baseline and optimized designs are fair. The full optimization problem includes variables to modify the twist distribution, cross-sectional shape, and planform of the wing, in addition to the structural sizing variables. The horizontal tail is also allowed to change in size to match planform changes in the wing.

As an intermediary step, we also optimize the wing without planform variables. To differentiate between them we designate the reduced problem S+T (shape and twist), while the full problem is labeled S+T+P (shape, twist, and planform). The result of the S+T problem informs our assessment of the success of the S+T+P optimization. While we fully expect the S+T+P optimized design to perform better than the baseline, it would be difficult to determine whether the improvements were due to the merits of MDO or to missing constraints in the optimization problem. The result of the S+T optimization provides a clearer picture of the benefit of simultaneously optimizing the structure, planform, and cross-sectional shape of the wing.

For the T+S+P problem, the tail size is constrained to maintain the tail volume coefficient of the baseline design. Including a tail scaling variable and a tail sizing constraint in the optimization problem is uncommon in the literature, and the effects are not obvious. The final case presented in this section is a control for the tail sizing experiment and does not allow the horizontal tail to scale. In all other respects it is identical to the T+S+P problem.

5.4.1 Tail Sizing Constraint

Adjustments to the wing planform can have global consequences on the design of the aircraft. For instance, we must take care that the planform of the wing does not change such that the tail is underpowered to counter the wing moment. This could occur due to changes in the location of the mean aerodynamic chord or the overall area of the wing. Raymer⁸¹ suggests the use of a tail volume coefficient to

match the moments produced by the tail and the wing. This parameter is purely geometrical, so it was fairly straightforward to add it as a function in pyConcept. The tail volume coefficient is defined as

$$c_{\text{tail}} = \frac{L_{\text{tail}}S_{\text{tail}}}{\bar{c}S_{\text{wing}}} \quad (5.4)$$

where

$$L_{\text{tail}} = \bar{x}_{\text{tail}} - \bar{x}_{\text{wing}}$$

and \bar{x} is the quarter-chord position of the mean aerodynamic chord (\bar{c}) of the lifting surface. Each of the parameters used in the expression for c_{tail} is updated in pyConcept based on changes in the FFDs of the wing and the tail. Additional geometric variables are assigned to the FFD enclosing the horizontal tail to allow it to scale up or down to match the changes in the wing planform.

5.4.2 Optimization Problem

The full optimization problem, corresponding to the S+T+P case, is listed in Table 5.6. The optimization problem is very similar to the problem solved in Section 5.3, so here we will only make note of the differences. The first major difference is that the objective function is simply the composite fuel burn, rather than a weighted average of fuel burn and OEW. The geometric parametrization is also slightly different for the R-jet. The twist distribution is controlled at six spanwise stations, whereas the chord can only scale at the mid-span kink, the wingtip, and the tip of the winglet. The tail rotates to trim the aircraft and scales to match changes in the planform (in accordance with the tail volume coefficient constraint). For this problem, we include a constraint to ensure a 30% margin to buffet in cruise flight. The buffet constraint is formulated based on the method of Kenway and Martins⁹⁵ and is applied at the nominal cruise flight condition under a 1.3g load. The tail volume coefficient is constrained to be greater than or equal to the baseline value.

5.4.3 Results

The results of the optimization are given in Table 5.7 in terms of percentage difference from the baseline. Fuel burn decreases across all missions, with the greatest percentage difference in fuel burn coming from the long-range mission. The S+T design reduces fuel burn by 1.8% compared to the baseline. With the addition of planform variables, the optimized design achieves an additional 1.8% reduction in the objective, double that of the S+T design. The structural weight of the S+T+P design is roughly 5% greater than the baseline, whereas the S+T design actually reduces the weight by slightly more than 8%. The increase in weight on the S+T+P wing is necessary to support an increased aspect ratio that is more than 25% greater

Table 5.6: R-jet optimization problem

		Quantity		
minimize	$\overline{W}_{\text{fuel}}$	1		
	with respect to	Angle of attack	5	
		Tail rotation	5	
		Twist	6	
		Span	1	
		Sweep	1	
		Chord scaling	3	
		Tail scaling	1	
		Sectional shape	272	
		Structural sizing	380	
		Panel lengths	146	
		Fuel load tractions	156	
		Total fuel mass	4	
		Mission fuel weights	12	
			Total number of design variables	992
		subject to	Range constraints	3
			Trim constraints ($L = W, C_{M,y} = 0$)	10
Structural failure constraints	4			
Planform area cannot decrease	1			
Tail volume coefficient constraint	1			
Wing geometric constraints	180			
Buffet onset constraint (1.3g)	1			
Miscellaneous mission constraints	159			
Structural sizing adjacency constraints	450			
Panel length consistency constraints	146			
Fuel traction consistency constraints	156			
Total fuel mass constraints	4			
Fuel volume constraints	4			
	Total number of design constraints	1119		

than the baseline design. Interestingly, the lift-to-drag ratio increases for both the nominal and long-range missions, but decreases for the high-speed mission. This is due to the fact that the longer range missions burn more fuel and therefore offer more in terms of fuel burn reduction than is lost on the high-speed mission.

Airfoil slices and C_p distributions of the baseline, S+T, and S+T+P designs are shown in Figure 5.9. The most noticeable change to the S+T airfoils is a decrease in thickness-to-chord ratio, especially on the outboard wing. The S+T+P airfoils feature a decrease in twist and more pronounced aft camber compared to the baseline. Both optimized designs have a more gradual pressure rise at the nominal design point, but at the high-speed flight condition, a shock is forming where none existed on the baseline wing. Remarkably, the C_p distributions for the long-range design point are fairly similar on the three designs.

Spanwise distributions of lift, twist, and t/c are plotted with respect to normalized span in Figure 5.10. The lift distribution for the S+T design is nearly identical to that of the baseline wing. Since it has a longer span, the S+T+P design can afford to generate less lift on the outboard wing, thereby reducing the bending moment caused by tip loads. In the plot of twist distributions, the effects of passive load al-

Table 5.7: Relative difference between baseline and optimized R-jet designs.

	S+T	S+T+P	S+T+P (fixed tail)
<i>Fuel burn</i>			
Nominal	-1.9	-3.6	-3.3
Long-range	-2.6	-4.8	-4.5
High-speed	-0.7	-1.9	-1.9
Combined	-1.8	-3.6	-3.4
<i>Mass</i>			
OEW	-1.1	+0.7	+0.5
Wing	-8.1	+4.9	+3.8
<i>L/D</i>			
Nominal	+3.1	+6.7	+5.9
Long-range	+3.6	+8.0	+7.2
High-speed	-1.9	-1.6	-1.4
<i>Geometry</i>			
Aspect ratio	0.0	+26.5	+28.8
Span (meters)	0.0	+12.5	+13.5
Sweep (degrees)	0.0	-24.2	-15.2




Figure 5.9: Comparison of airfoils and C_p distributions for different cases.

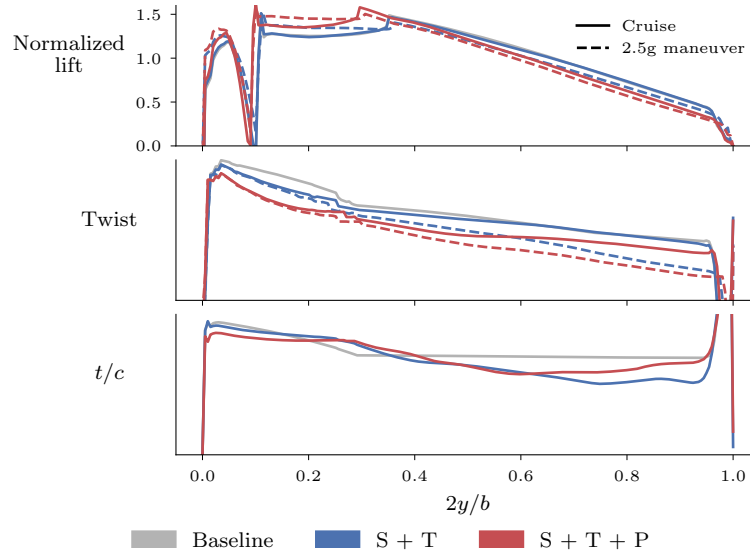


Figure 5.10: Comparison of spanwise lift, twist, and t/c distributions for different cases.

leviation at the maneuver condition are readily apparent. The bend-twist coupling in the structure produces a moment to untwist the wing when deflected under the heavy 2.5g load. As observed in the airfoil slices, the plot of t/c shows substantial reduction in thickness in the outboard wing on both optimized designs.

The most striking differences between the baseline and S+T+P designs are the increase in aspect ratio and decrease in wing sweep, as shown in Figure 5.11. Basic aerodynamic theory tells us that induced drag and wave drag vary inversely with span and sweep, respectively. However, both span and sweep are tightly coupled with the structural response of the wing as well. Any increase in span or sweep generally requires a heavier wing structure to support the increased moment arm. To determine the correct tradeoff between span and sweep, it is critical to model the wing aerostructurally. In this case, the optimizer increases span to reduce induced drag at the cost of a heavier wing. At the same time, the decrease in wing sweep lightens the burden on the wing while possibly degrading the high-speed performance of the wing (as evidenced by the shocks shown on the high-speed design point in Figure 5.9). The right side of Figure 5.11 shows the deflected state of the optimized wings at cruise and maneuver conditions.

When the scaling of the horizontal tail is an active design variable, the optimizer chooses to shrink it in order to reduce drag. Since the wing area cannot decrease, Equation 5.4 dictates that the tail area can only decrease if the wing sweep is reduced. With the horizontal tail area fixed, the optimizer still reduces the wing sweep, but not as much as in the former case. In the case without fixed tail area, the optimizer is actually degrading the performance of the wing slightly in order to reap the benefits of a smaller horizontal tail which leads to better overall performance. Whereas the




Figure 5.11: Views of planform and deflected wing shape for the optimized wings. The S+T wing planform is identical to the baseline wing.

designer might be reluctant to sacrifice wing performance and therefore arrive at a suboptimal overall design, the optimizer has no such inhibitions and is able to find proper tradeoff between wing performance and tail size. The difference in the sizes of the horizontal tails can be seen in Figure 5.11.

Contours of the structural failure and buckling criteria as well as the skin thickness and aeroelastic deflection for the baseline and S+T+P designs are shown in Figure 5.12. The panel thickness on both the upper and lower skins increases to support the longer wingspan.

5.5 Summary

In Section 5.3, we demonstrated that different methods of calculating fuel burn can impact the optimized design. We experimented with a method for mission analysis that combined low-fidelity aerodynamic analyses along the climb and descent segments with a high-fidelity aerostructural analysis to represent the mid-cruise point. With this hybrid method, the fuel burn on the climb and descent segments is included in the objective function and is dependent on the geometric parametrization. For short-haul missions, where the aircraft burns a significant portion of the total fuel burn in climb, the correct tradeoffs are more likely to be achieved with this approach.

In Section 5.4, we test the merits of the MACH framework on the optimization of an Embraer regional jet. The wing optimized with shape, twist, and planform variables burns 3.6% less fuel than the baseline design and 1.8% less fuel than a wing optimized with only shape and twist variables. When all the variables are included, the wing span increases substantially and the wing sweep is reduced compared to the baseline. The significant differences between the baseline and optimized wings showcase the advantage of simultaneously optimizing all variables while

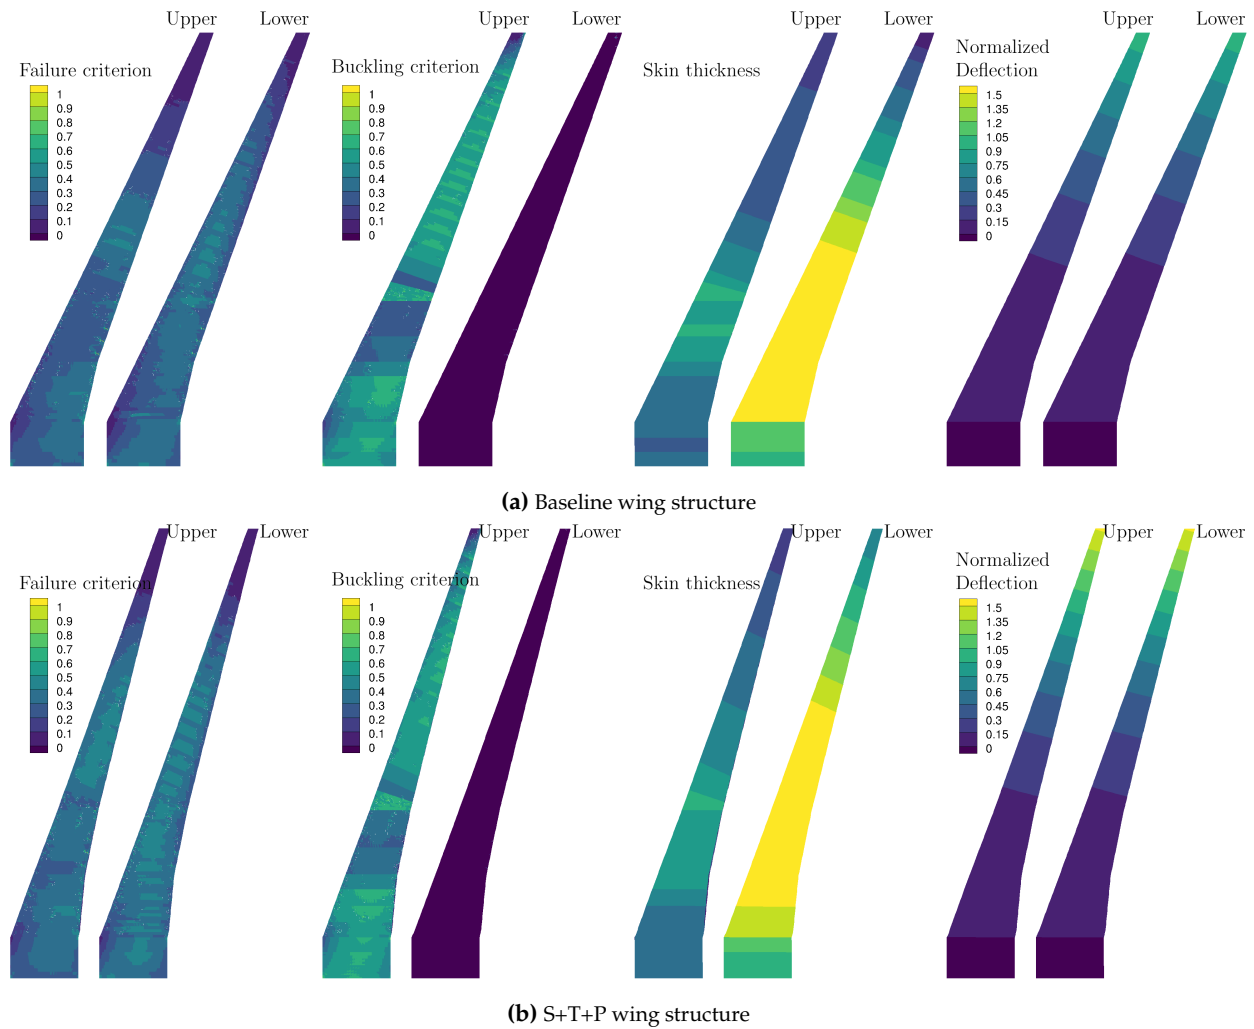


Figure 5.12: Comparison of baseline and optimized wingbox structures. The wing deflection is normalized by the maximum tip deflection on the baseline wing.

considering the aerostructural coupling of the wing. As mentioned previously, both span and sweep are tightly coupled with the structural response of the wing as well. To determine the correct tradeoff between span and sweep, it is critical to model the wing aerostructurally.

It is likely that in reality, the wing design is influenced by constraints and requirements that are neglected in this study. Similarly, it is possible that the level of detail in the modeling is not sufficient to yield a realistic result, and that adding more detail would change the result. Even so, these results give every indication that the simultaneous design of wingbox structure, wing planform, and wing cross-sectional shapes will yield a different wing design than when each is designed separately.

Chapter 6

Final Remarks

The goal of this dissertation is to demonstrate the versatility and usefulness of high-fidelity MDO in aircraft wing design. Broadly speaking, I sought to achieve this goal by solving a variety of wing design optimization problems and then carefully analyzing the behavior and output of the optimizer. This dissertation is especially targeted toward demonstrating the ability of a gradient-based optimizer to explore the design and balance multidisciplinary tradeoffs effectively. I also made special efforts to address common concerns and provide practical suggestions to the designer. Section 6.1 contains a summary of specific key results and conclusions from the dissertation. This is followed by enumerated contributions in Section 6.2 and potential avenues for future research in Section 6.3.

6.1 Conclusions

In Chapter 3, I explored the design space of a subsonic wing benchmark problem with the aim of demystifying multimodality. I found that multimodality is highly dependent on the tradeoff between skin friction drag and induced drag. For the drag minimization problem with only chord variables I showed that the number of local minima goes to infinity as lift goes to zero. As lift and induced drag increase, the optimizer favors an elliptical chord distribution that will maximize span efficiency and the design space becomes unimodal. Therefore it is not only the optimal chord distribution that is dependent on the skin-friction/induced drag tradeoff, but the topology of the design space itself. If the design space is a collection of peaks and valleys, varying the ratio of skin-friction to induced drag does more than just scale the peaks and valleys, rather, it introduces an entirely new landscape. Conversely, multimodality in the span-constrained, nonplanar wing design problem varies inversely with the ratio of skin-friction to induced drag. For a high lift condition, the optimizer will produce either an upturned or a downturned winglet; both

offer similar reductions in induced drag. As C_L decreases and skin-friction drag overtakes induced drag, the wing is flattened to reduce skin-friction drag and the design space becomes unimodal. These insights are important because they improve the designer's understanding of optimizer behavior. Rather than dismiss optimization as unsuitable for industrial work upon seeing a strange result or finding multiple answers to the same problem, the designer can reformulate the problem to obtain a more satisfactory result. An example of a fix that was suggested for the multimodal chord distribution is a constraint to force a monotonic or linear chord distribution.

The primary purpose of Chapter 4 was to demonstrate that MDO is an effective means of exploring the aerostructural wing design space. I showed that an optimizer could start from a rectangular, constant cross-section wing and traverse the design space to arrive at a conventional transonic swept wing with custom airfoils and optimally sized structure. Moreover, the optimizer arrived at the same design (within a small tolerance) when the optimization began from a swept planform. This demonstration inspires confidence that a gradient-based optimizer is not doomed to converge to suboptimal local minima when applied to an aerostructural optimization problem. To reduce the computation cost of this study, I used a multi-level optimization process, in which the CFD grids are refined in successive optimization runs, to reduce the overall computational cost by 40%.

Chapter 4 also introduces a novel method to improve off-design robustness in optimized wings. First I showed that both single-point and multipoint cruise-optimized designs exhibit massive separation at a low-speed, high-lift flight condition representative of a typical climb profile. A separation constraint applied at the climb condition restores healthy flow without severely degrading cruise performance. Typically in wing design optimization, leading edge thickness constraints are used to prevent the optimizer from excessively reducing the leading edge radius. A round leading edge prevents separation at large incidence angles, but is not undesirable for optimum cruise performance. With the proposed low-speed separation constraint, adequate curvature was preserved on the leading edge without having to resort to thickness constraints.

In Chapter 5, I investigated the redesign of an Embraer regional jet. For a regional jet, the typical mission is no longer dominated by the cruise segment, so it is important to capture the performance of the aircraft in climb and descent. I experimented with various methods of aggregating fuel burn in the objective function of the optimization problem. The results indicate that the planform of the optimized varies significantly based on whether climb and descent fuel burn is included in the objective. More precisely, the optimal aspect ratio depends on the ratio between the variable fuel burn and the variable structural weight. The

optimizer can reduce fuel burn by increasing aspect ratio, which improves L/D at the expense of increased structural weight. Conversely, the optimizer can reduce fuel burn by decreasing aspect ratio, which reduces structural weight but worsens L/D . For a regional jet, when fuel burn is computed across the entire mission, L/D has a larger relative impact on total fuel burn than structural weight, so the optimal aspect ratio is larger. If the fuel burn is calculated solely on the cruise segment, L/D has less influence on the final design and the optimal aspect ratio is less. This study demonstrates the effectiveness of using MDO to determine the correct tradeoff between competing multidisciplinary engineering systems. It also is a cautionary example of the importance of providing the optimizer with the correct information so that it can balance the tradeoff properly.

The final section of Chapter 5 illustrates the benefits of simultaneously optimizing twist, airfoil shape, planform, and structural members in a wing design problem. Compared to the baseline design, the optimized wing burns 3.6% less fuel burn than the baseline design. The optimized wing features reduced sweep, longer span, and increased structural weight. I also compared the result of this full optimization problem to one in which the planform is restricted to the initial design and found a 1.8% reduction in fuel burn for the former case. Without high-fidelity analysis in the optimization loop, it would not have been possible to properly model the tradeoffs between all of these design parameters. This is not to say that low-fidelity models do not have a place in the design process—on the contrary, low-fidelity models were used many times in this dissertation to explore specific cases. All models are approximations of reality, and the usefulness of a given model at “predicting the incremental effect of a proposed design change”¹⁰ must be determined based on its assumptions and shortcomings. The closing argument of this dissertation is that high-fidelity, multidisciplinary models are required for detailed design of aircraft wings, and that gradient-based optimization is a viable tool for exploring the design space produced by these models.

6.2 Contributions

Most of the tools used in this dissertation had been developed previously. I contributed updates to the MACH framework in the course of the research, most conspicuously to the pyGeo and MissionAnalysis codes. However, the major original contributions in this dissertation are more concerned with application and analysis than implementation. The most significant of these are listed here.

1. *Developed new insights into the tradeoff between skin-friction and induced drag and its impact on multimodality and optimization.* The understanding that there is a tradeoff between skin-friction and induced drag is not revolutionary. As

explained in Section 3.2.2, this result is based on fundamental aerodynamic principles. However, the recognition that this tradeoff has an important impact on multimodality in the design space is a novel contribution. I demonstrated this effect in a chord distribution optimization and a nonplanar wing optimization.

2. *Demonstrated practical ways to reduce multimodality in the design space.* Several papers have been written reporting significant multimodality in the wing planform design problem, however, there is little explanation of the fundamental reasons for this multimodality or how to deal with it. I showed that in many cases, multimodality is dependent on the physics model and flow condition of the design point, both of which are under the control of the designer. In cases where the design point inherently promotes a multimodal design space, I demonstrated the use of geometric constraints to remove the multimodality.
3. *Demonstrated exploratory high-fidelity aerostructural optimization by converging from a plank wing to a conventional transonic wing.* The transformation from a constant cross-section, rectangular wing to a typical, swept transonic wing is the most extreme example of aerostructural design space exploration in the literature. This result is valuable because it demonstrates that a gradient-based optimizer is not destined to get stuck at a local minimum close to the initial starting point. On the contrary, the optimizer is able to traverse the design space to arrive at a completely different wing design. While this has been shown in ASO problems, aerostructural optimization problems are generally more difficult to converge and cannot be assumed to exhibit the same behavior as purely aerodynamic optimization problems.
4. *Pioneered low-speed separation constraint to improve off-design robustness.* One of the major limitations of high-fidelity wing design optimization is the difficulty of simultaneously maximizing performance at a cruise condition while preserving flightworthiness at low-speed, high-lift conditions. If the off-design performance is considered in the problem formulation, it is usually with a low-fidelity model, and if not, geometric constraints must be put in place to prevent the optimizer from exploiting the omitted requirements. I pioneered the use of a separation constraint at the low-speed, high-lift flight condition to improve off-design robustness. The new constraint induced significant changes in the optimized wing shape, most notably the restoration of a suitably rounded leading edge. As stated in the text, the separation constraint was previously developed by Kenway and Martins⁹⁵ to prevent transonic buf-

fet. The contribution in this dissertation is the application of the separation constraint to preserve performance at low-speed, high-lift conditions.

5. *Investigated impact of mission analysis method on optimization of short-haul aircraft.* In the literature, fuel burn is commonly approximated using the Bréguet range formula to simplify the optimization problem. I showed that, for a regional jet, the method of computing fuel burn has a significant impact on the optimized wing design, especially the aspect ratio. I demonstrated the use of a novel hybrid mission analysis, in which the climb and descent segments are represented by a low-fidelity conceptual aerodynamics model and the cruise segment is represented by the Bréguet range equation with a high-fidelity analysis point.
6. *Demonstrated industrial application of MDO on Embraer regional jet.* Industrial applications are scarce in the literature, giving the impression that MDO is not yet ready for such problems. In this collaboration with Embraer, I used MDO to optimize the wing of an authentic regional jet. The optimized wing is significantly different than the baseline design, showcasing the optimizer's facility for balancing multidisciplinary tradeoffs in a way is not possible for discrete engineering teams.

6.3 Recommendations for Future Work

1. *Improve convergence of aerostructural optimizations.* In general, aerostructural optimizations are more difficult to converge and do not converge as tightly as aerodynamic-only optimizations. This is, at least in part, due to the high condition number of the stiffness matrix for shell finite elements, which limits the achievable numerical precision for both structural and aerostructural analyses and gradient computations. Additionally, the KS functions used to aggregate structural failure criteria are highly nonlinear and can create problems for gradient-based optimizers. Experience suggests that convergence deteriorates as the the complexity (i.e. size) of the optimization problem increases. A more rigorous study of aerostructural optimization problems should be conducted to determine the extent to which of each of these factors affects optimization convergence. It would also be an important contribution to establish more general expectations for the convergence behavior of aerostructural optimization problems.
2. *Allow for changes in the wing-fuselage intersection.* In this dissertation, the geometric parametrizations did not allow the wing-fuselage intersection to change shape. Allowing the intersection region to change would open up

new possibilities in terms of the tradeoffs that could be considered in the optimization problem. However, intersections are notoriously troublesome for gradient-based optimization. Work is currently being done in the MDO Lab to address these issues.

3. *Explore more general structural parametrization schemes, both in wingbox layout and material properties.* For the aerostructural problems in this dissertation, the structure was parametrized with local sizing parameters—both the layout and material properties of the ribs and spars were fixed. A more flexible parametrization of either or both of the layout and material properties could lead to greater coupling between the OML and wingbox designs. This has been shown in previous work with regard to tow-steered composite fibers, which could be oriented to more efficiently counteract aerodynamic loads. Other work has looked into structural optimization of the location, orientation, and curvature of ribs and spars in a wingbox layout, but this has not been extended to high-fidelity aerostructural optimization.
4. *Verify low-speed, high-lift separation constraint.* While the separation constraint at a low-speed, high-lift condition demonstrably improved off-design performance in the results presented herein, further work is needed to correlate the constraint with industry or regulatory requirements.
5. *Conduct experimental validation of optimized designs.* Some of the designs produced by the optimizer in this dissertation are quite nonintuitive. In the literature, experimental validation is rarely used to substantiate the products of optimization, both due to the cost and the fact that most researchers in the field of MDO do purely computational work. In my opinion, experimental validation of the results of numerical wing optimization would be a compelling contribution, especially in the case of nonintuitive designs that deviate significantly from conventional wings.

Bibliography

[Citing pages are listed after each reference.]

- 1 John D. Anderson, Jr. *Introduction to Flight*. McGraw-Hill, 7th edition, 2012. [p. 2]
- 2 John D. Anderson, Jr. *Aircraft Performance and Design*. McGraw-Hill, 1999. [p. 2]
- 3 Ludwig Prandtl. Über flüssigkeitsbewegung bei sehr kleiner reibung. In *Proceedings of the 3rd International Mathematical Congress*, pages 484–491, Heidelberg, 1904. [p. 2]
- 4 Ludwig Prandtl. Theory of lifting surfaces. Technical Report 9, National Advisory Committee for Aeronautics, July 1920. [p. 2]
- 5 Max M. Munk. General theory of thin wing sections. Technical Report 142, National Advisory Committee for Aeronautics, 1922. [p. 2]
- 6 Ira H. Abbott. Airfoils - significance and early development. In *The evolution of aircraft wing design; Proceedings of the Symposium*. American Institute of Aeronautics and Astronautics, March 1980. doi:[10.2514/6.1980-3033](https://doi.org/10.2514/6.1980-3033). [pp. 2 and 3]
- 7 Theodore Theodorsen. Theory of wing sections of arbitrary shape. Technical Report 411, National Advisory Committee for Aeronautics, 1931. [p. 2]
- 8 A. M. O. Smith and J. Pierce. Exact solution of the neumann problem. calculation of non-circulatory plane and axially symmetric flows about or within arbitrary boundaries. Technical Report ES 26988, Douglas Aircraft Company, Long Beach, CA, April 1958. [p. 3]
- 9 John L. Hess and A. M. O. Smith. Calculation of potential flow about arbitrary bodies. *Progress in Aerospace Sciences*, 8(1–138), 1967. doi:[10.1016/0376-0421\(67\)90003-6](https://doi.org/10.1016/0376-0421(67)90003-6). [p. 3]

- 10 John L. Hess. Panel methods in computational fluid dynamics. *Annual Review of Fluid Mechanics*, 22(1):255–274, January 1990. [pp. 3 and 98]
- 11 W. Mangler. Die berechnung eines tragflügelprofiles mit vorgeschriebener druckverteilung. *Jahrbuch der deutschen Luftfahrtforschung*, 1:46–53, 1938. [p. 3]
- 12 M. James Lighthill. A new method of two-dimensional aerodynamic design. Technical Report 2112, Aeronautical Research Council, 1945. [p. 3]
- 13 Lucien A. Schmit. Structural Design by Systematic Synthesis. In *2nd Conference on Electronic Computation*, pages 105–132, New York, NY, 1960. ASCE. [p. 4]
- 14 Raymond M. Hicks, Earll M. Murman, and Garret N. Vanderplaats. An assessment of airfoil design by numerical optimization. Technical Report NASA TM X-3092, NASA Ames Research Center, Moffet Field, CA, July 1974. [p. 4]
- 15 Raymond M. Hicks and Preston A. Henne. Wing design by numerical optimization. *Journal of Aircraft*, 15(7):407–412, 1978. [pp. 4, 5, and 7]
- 16 Olivier Pironneau. On optimum profiles in Stokes flow. *Journal of Fluid Mechanics*, 59(01):117–128, 1973. doi:[10.1017/S002211207300145X](https://doi.org/10.1017/S002211207300145X). [p. 5]
- 17 Olivier Pironneau. On optimum design in fluid mechanics. *Journal of Fluid Mechanics*, 64(1):97–110, 1974. doi:[10.1017/S0022112074002023](https://doi.org/10.1017/S0022112074002023). [p. 5]
- 18 Antony Jameson. Aerodynamic design via control theory. *Journal of Scientific Computing*, 3(3):233–260, September 1988. doi:[10.1007/BF01061285](https://doi.org/10.1007/BF01061285). [p. 5]
- 19 Antony Jameson. Automatic Design of Transonic Airfoils to Reduce the Shock Induced Pressure Drag. In *Proceedings of the 31st Israel Annual Conference on Aviation and Aeronautics*, pages 5–17, Technion-Israel, Haifa, Israel, 1990. [p. 5]
- 20 Antony Jameson. Optimum aerodynamic design via boundary control. Technical Report NASA-CR-195882, NASA, August 1994. [p. 5]
- 21 Jamshid A. Samareh. Survey of shape parameterization techniques for high-fidelity multidisciplinary shape optimization. *AIAA Journal*, 39(5):877–884, May 2001. doi:[10.2514/2.1391](https://doi.org/10.2514/2.1391). [p. 5]
- 22 James J. Reuther, Antony Jameson, J. Farmer, Luigi Martinelli, and David Saunders A. Aerodynamic shape optimization of complex aircraft configurations via an adjoint formulation. In *Proceedings of the 34th AIAA Aerospace Sciences Meeting and Exhibit*, Reno, Nevada, January 1996. AIAA 1996-0094. [p. 5]

- 23 James J. Reuther, Antony Jameson, Juan J. Alonso, Mark J. Rimlinger, and David A. Saunders. Constrained multipoint aerodynamic shape optimization using an adjoint formulation and parallel computers, part 1. *Journal of Aircraft*, 36(1):51–60, Jan 1999. [p. 6]
- 24 James J. Reuther, Antony Jameson, Juan J. Alonso, Mark J. Rimlinger, and David A. Saunders. Constrained multipoint aerodynamic shape optimization using an adjoint formulation and parallel computers, part 2. *Journal of Aircraft*, 36(1):61–74, Jan 1999. [pp. 6 and 8]
- 25 Antony Jameson, Luigi Martinelli, and Niles A. Pierce. Optimum aerodynamic design using the Navier–Stokes equations. *Theoretical and Computational Fluid Dynamics*, 10(1–4):213–237, 1998. doi:10.1007/s001620050060. [p. 6]
- 26 W. K. Anderson and V. Venkatakrishnan. Aerodynamic design optimization on unstructured grids with a continuous adjoint formulation. *Computers and Fluids*, 28(4):443–480, 1999. doi:10.1016/S0045-7930(98)00041-3.
- 27 Eric J. Nielsen and W. Kyle Anderson. Aerodynamic design optimization on unstructured meshes using the Navier–Stokes equations. *AIAA Journal*, 37(11): 1411–1419, 1999. doi:10.2514/2.640. [p. 6]
- 28 Zhoujie Lyu, Gaetan K. W. Kenway, Cody Paige, and Joaquim R. R. A. Martins. Automatic differentiation adjoint of the Reynolds-averaged Navier–Stokes equations with a turbulence model. In *21st AIAA Computational Fluid Dynamics Conference*, San Diego, CA, Jul. 2013. [pp. 6 and 21]
- 29 Lana Osusky, Howard Buckley, Thomas Reist, and David W. Zingg. Drag minimization based on the navier–stokes equations using a newton–krylov approach. *AIAA Journal*, 53(6):1555–1577, Jun 2015. [p. 6]
- 30 Joël Brezillon and Richard P. Dwight. Applications of a discrete viscous adjoint method for aerodynamic shape optimisation of 3D configurations. *CEAS Aeronautical Journal*, 3(1):25–34, 2012. doi:10.1007/s13272-011-0038-0. [p. 6]
- 31 Song Chen, Zhoujie Lyu, Gaetan K. W. Kenway, and Joaquim R. R. A. Martins. Aerodynamic shape optimization of common research model wing–body–tail configuration. *Journal of Aircraft*, 53(1):276–293, January 2016.
- 32 Thomas A. Reist and David W. Zingg. High-fidelity aerodynamic shape optimization of a lifting-fuselage concept for regional aircraft. *Journal of Aircraft*, 54(3):1085–1097, May 2017. [p. 6]

- 33 Zhoujie Lyu and Joaquim R. R. A. Martins. Aerodynamic design optimization studies of a blended-wing-body aircraft. *Journal of Aircraft*, 51(5):1604–1617, Sep 2014. doi:[10.2514/1.C032491](https://doi.org/10.2514/1.C032491). [p. 6]
- 34 Zhoujie Lyu and Joaquim R. R. A. Martins. Aerodynamic shape optimization of an adaptive morphing trailing edge wing. *Journal of Aircraft*, 52(6):1951–1970, November 2015. doi:[10.2514/1.C033116](https://doi.org/10.2514/1.C033116).
- 35 Ney R. Secco and Joaquim R. R. A. Martins. RANS-based aerodynamic shape optimization of a strut-braced wing with overset meshes. *Journal of Aircraft*, 56(1):217–227, January 2019. doi:[10.2514/1.C034934](https://doi.org/10.2514/1.C034934). [p. 6]
- 36 Jason E. Hicken and David W. Zingg. Induced-drag minimization of nonplanar geometries based on the euler equations. *AIAA Journal*, 48(11):2564–2575, 2010. doi:[10.2514/1.J050379](https://doi.org/10.2514/1.J050379). [p. 6]
- 37 Hugo Gagnon and David W. Zingg. Euler-equation-based drag minimization of unconventional aircraft configurations. *Journal of Aircraft*, 53(5):1361–1371, September 2016. doi:[10.2514/1.C033591](https://doi.org/10.2514/1.C033591). [p. 6]
- 38 Andrew Ning and Ilan Kroo. Multidisciplinary considerations in the design of wings and wing tip devices. *Journal of Aircraft*, 47(2):534–543, Mar 2010. doi:[10.2514/1.41833](https://doi.org/10.2514/1.41833). [pp. 6 and 66]
- 39 Hugo Gagnon and David W. Zingg. Two-level free-form and axial deformation for exploratory aerodynamic shape optimization. *AIAA Journal*, 53(7):2015–2026, 2015. doi:[10.2514/1.J053575](https://doi.org/10.2514/1.J053575). [p. 6]
- 40 Antony Jameson and Kui Ou. 50 years of transonic aircraft design. *Progress in Aerospace Sciences*, 47(5):308 – 318, 2011. doi:[10.1016/j.paerosci.2011.01.001](https://doi.org/10.1016/j.paerosci.2011.01.001). [p. 6]
- 41 S.N. Skinner and H. Zare-Behtash. State-of-the-art in aerodynamic shape optimisation methods. *Applied Soft Computing*, 62:933–962, January 2018. doi:[10.1016/j.asoc.2017.09.030](https://doi.org/10.1016/j.asoc.2017.09.030). [p. 6]
- 42 Raphael T. Haftka. Automated procedure for design of wing structures to satisfy strength and flutter requirements. Technical Report TN D-7264, NASA, July 1973. [p. 6]
- 43 Raphael T. Haftka. Optimization of flexible wing structures subject to strength and induced drag constraints. *AIAA Journal*, 15(8):1101–1106, Aug 1977. [p. 6]
- 44 Dietrich Küchemann. *The Aerodynamic Design of Aircraft*. Pergamon Press, Oxford, UK, 1st edition, 1978. [p. 6]

- 45 Joaquim R. R. A. Martins and Andrew B. Lambe. Multidisciplinary design optimization: A survey of architectures. *AIAA Journal*, 51(9):2049–2075, September 2013. [p. 6]
- 46 N. Maman and Charbel Farhat. Matching fluid and structure meshes for aeroelastic computations: A parallel approach. *Comput. and Struc.*, 54:779–785, 1995. [p. 6]
- 47 S. A. Brown. Displacement extrapolations for CFD+CSM aeroelastic analysis. In *38th Structures, Structural Dynamics, and Materials Conference*, April 1997. doi:10.2514/6.1997-1090. [p. 16]
- 48 Charbel Farhat, M. Lesoinne, and P. Le Tallec. Load and motion transfer algorithms for fluid/structure interaction problems with non-matching discrete interfaces: Momentum and energy conservation, optimal discretization and application to aeroelasticity. *Computer Methods in Applied Mechanics and Engineering*, 157(1):95 – 114, 1998. doi:10.1016/S0045-7825(97)00216-8. [p. 6]
- 49 K. Maute, M. Nikbay, and Charbel Farhat. Coupled analytical sensitivity analysis and optimization of three-dimensional nonlinear aeroelastic systems. *AIAA Journal*, 39(11):2051–2061, 2001. [p. 6]
- 50 Joaquim R. R. A. Martins, Juan J. Alonso, and James J. Reuther. A coupled-adjoint sensitivity analysis method for high-fidelity aero-structural design. *Optimization and Engineering*, 6(1):33–62, March 2005. [p. 7]
- 51 Joaquim R. R. A. Martins, Juan J. Alonso, and James J. Reuther. High-fidelity aerostructural design optimization of a supersonic business jet. *Journal of Aircraft*, 41(3):523–530, May 2004. [p. 7]
- 52 Gaetan K. W. Kenway and Joaquim R. R. A. Martins. Multipoint high-fidelity aerostructural optimization of a transport aircraft configuration. *Journal of Aircraft*, 51(1):144–160, Jan 2014. [pp. 7 and 66]
- 53 David A. Burdette and Joaquim R. R. A. Martins. Design of a transonic wing with an adaptive morphing trailing edge via aerostructural optimization. *Aerospace Science and Technology*, 81:192–203, October 2018. doi:10.1016/j.ast.2018.08.004.
- 54 Timothy R. Brooks, Gaetan K. W. Kenway, and Joaquim R. R. A. Martins. Benchmark aerostructural models for the study of transonic aircraft wings. *AIAA Journal*, 56(7):2840–2855, July 2018. doi:10.2514/1.j056603. [pp. 48, 66, 67, and 74]

- 55 Shahriar Khosravi and David W. Zingg. Aerostructural perspective on winglets. *Journal of Aircraft*, 2017. [pp. 34 and 67]
- 56 Shahriar Khosravi and David W. Zingg. Aerostructural optimization of drooped wings. *Journal of Aircraft*, 55(3):1261–1268, May 2018. doi:10.2514/1.C034605. [p. 7]
- 57 Oleg Chernukhin and David W. Zingg. Multimodality and global optimization in aerodynamic design. *AIAA journal*, 51(6):1342–1354, 2013. [p. 7]
- 58 Zhoujie Lyu, Gaetan K. W. Kenway, and Joaquim R. R. A. Martins. Aerodynamic shape optimization investigations of the common research model wing benchmark. *AIAA Journal*, 53(4):968–985, April 2015. [pp. 7, 8, 48, 54, 60, and 66]
- 59 David Koo and David W. Zingg. Investigation into aerodynamic shape optimization of planar and nonplanar wings. *AIAA Journal*, 56(1):250–263, Jan 2018. doi:10.2514/1.J055978.
- 60 Yin Yu, Zhoujie Lyu, Zelu Xu, and Joaquim R. R. A. Martins. On the influence of optimization algorithm and starting design on wing aerodynamic shape optimization. *Aerospace Science and Technology*, 75:183–199, April 2018. doi:10.1016/j.ast.2018.01.016. [pp. 7 and 12]
- 61 Gregg M. Streuber and David W. Zingg. Investigation of multimodality in aerodynamic shape optimization based on the reynolds averaged navier-stokes equations. In *35th AIAA Applied Aerodynamics Conference*, Denver, CO, Jun 2017. American Institute of Aeronautics and Astronautics. [pp. 7 and 20]
- 62 Gregg Streuber and David W. Zingg. A parametric study of multimodality in aerodynamic shape optimization of wings. In *Multidisciplinary Analysis and Optimization Conference*, Atlanta, GA, Jun 2018. American Institute of Aeronautics and Astronautics. [p. 7]
- 63 Mark Drela. Pros and cons of airfoil optimization. In D. A. Caughey and M. M. Hafez, editors, *Frontiers of Computational Fluid Dynamics*. World Scientific, 1998. [pp. 8, 62, 66, and 75]
- 64 Gaetan K. W. Kenway and Joaquim R. R. A. Martins. Multipoint aerodynamic shape optimization investigations of the common research model wing. *AIAA Journal*, 54(1):113–128, Jan 2016. doi:10.2514/1.J054154. [pp. 8, 66, and 67]
- 65 Gaetan K. W. Kenway, Graeme J. Kennedy, and Joaquim R. R. A. Martins. Scalable parallel approach for high-fidelity steady-state aeroelastic analysis and adjoint derivative computations. *AIAA Journal*, 52(5):935–951, May 2014. [pp. 10 and 16]

- 66 Andrew B. Lambe and Joaquim R. R. A. Martins. Extensions to the design structure matrix for the description of multidisciplinary design, analysis, and optimization processes. *Structural and Multidisciplinary Optimization*, 46:273–284, August 2012. doi:[10.1007/s00158-012-0763-y](https://doi.org/10.1007/s00158-012-0763-y). [p. 10]
- 67 Philip E. Gill, Walter Murray, and Michael A. Saunders. SNOPT: An SQP algorithm for large-scale constrained optimization. *SIAM Review*, 47(1):99–131, 2005. doi:[10.1137/S0036144504446096](https://doi.org/10.1137/S0036144504446096). [p. 12]
- 68 Thomas W. Sederberg and Scott R. Parry. Free-form deformation of solid geometric models. *SIGGRAPH Comput. Graph.*, 20(4):151–160, August 1986. doi:[10.1145/15886.15903](https://doi.org/10.1145/15886.15903). [p. 12]
- 69 Gaetan K. W. Kenway, Graeme J. Kennedy, and Joaquim R. R. A. Martins. A CAD-free approach to high-fidelity aerostructural optimization. In *13th AIAA/ISSMO Multidisciplinary Analysis Optimization Conference*, September 2010. doi:[10.2514/6.2010-9231](https://doi.org/10.2514/6.2010-9231). [p. 12]
- 70 Anil Yildirim, Gaetan K. W. Kenway, Charles A. Mader, and Joaquim R. R. A. Martins. A Jacobian-free approximate Newton–Krylov startup strategy for RANS simulations. *Journal of Computational Physics*, 397:108741, November 2019. doi:[10.1016/j.jcp.2019.06.018](https://doi.org/10.1016/j.jcp.2019.06.018). [p. 14]
- 71 Charles A. Mader, Joaquim R. R. A. Martins, Juan J. Alonso, and Edwin van der Weide. ADjoint: An approach for the rapid development of discrete adjoint solvers. *AIAA Journal*, 46(4):863–873, April 2008. doi:[10.2514/1.29123](https://doi.org/10.2514/1.29123). [p. 14]
- 72 Gaetan K. W. Kenway, Charles A. Mader, Ping He, and Joaquim R. R. A. Martins. Effective adjoint approaches for computational fluid dynamics. *Progress in Aerospace Sciences*, 110:100542, October 2019. doi:[10.1016/j.paerosci.2019.05.002](https://doi.org/10.1016/j.paerosci.2019.05.002). [p. 14]
- 73 Ney R. Secco, Gaetan K. W. Kenway, Ping He, Charles A. Mader, and Joaquim R. R. A. Martins. Efficient mesh generation and deformation for aerodynamic shape optimization. *AIAA Journal*, 2020. (Submitted). [p. 15]
- 74 Graeme J. Kennedy and Joaquim R. R. A. Martins. A parallel finite-element framework for large-scale gradient-based design optimization of high-performance structures. *Finite Elements in Analysis and Design*, 87:56–73, September 2014. doi:[10.1016/j.finel.2014.04.011](https://doi.org/10.1016/j.finel.2014.04.011). [p. 15]
- 75 Graeme J. Kennedy and Joaquim R. R. A. Martins. Parallel solution methods for aerostructural analysis and design optimization. In *Proceedings of*

- the 13th AIAA/ISSMO Multidisciplinary Analysis Optimization Conference*, Fort Worth, TX, September 2010. AIAA 2010–9308. [p. 16]
- 76 Graeme J. Kennedy and Joaquim R. R. A. Martins. A parallel aerostructural optimization framework for aircraft design studies. *Structural and Multidisciplinary Optimization*, 50(6):1079–1101, Aug 2014. doi:[10.1007/s00158-014-1108-9](https://doi.org/10.1007/s00158-014-1108-9). [p. 16]
- 77 Mustafa Cavcar. Bréguet range equation? *Journal of Aircraft*, 43(5):1542–1544, 2006. doi:[10.2514/1.17696](https://doi.org/10.2514/1.17696). [p. 17]
- 78 René Devillers. *La dynamique de l'avion*. Librairie aéronautique, 1918. [p. 17]
- 79 Joseph George Coffin. A Study of Airplane Ranges and Useful Loads. Technical Report 69, National Advisory Committee for Aeronautics, 1920. [p. 17]
- 80 Rhea P. Liem, Charles A. Mader, Edmund Lee, and Joaquim R. R. A. Martins. Aerostructural design optimization of a 100-passenger regional jet with surrogate-based mission analysis. *2013 Aviation Technology, Integration, and Operations Conference*, Aug 2013. doi:[10.2514/6.2013-4372](https://doi.org/10.2514/6.2013-4372). [pp. 17, 76, 78, and 81]
- 81 Daniel P. Raymer. *Aircraft Design: A Conceptual Approach*. AIAA, 5th edition, 2012. [pp. 18 and 89]
- 82 Daniel J. Poole, Christian B. Allen, and Thomas C. S. Rendall. Global optimization of wing aerodynamic optimization case exhibiting multimodality. *Journal of Aircraft*, 55(4):1576–1591, Jul 2018. doi:[10.2514/1.c034718](https://doi.org/10.2514/1.c034718). [p. 20]
- 83 Gaetan K. W. Kenway, Ney Secco, Joaquim R. R. A. Martins, Asitav Mishra, and Karthik Duraisamy. An efficient parallel overset method for aerodynamic shape optimization. In *Proceedings of the 58th AIAA/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference, AIAA SciTech Forum*, Grapevine, TX, January 2017. [p. 21]
- 84 John P. Jasa, John T. Hwang, and Joaquim R. R. A. Martins. Open-source coupled aerostructural optimization using Python. *Structural and Multidisciplinary Optimization*, 57(4):1815–1827, April 2018. doi:[10.1007/s00158-018-1912-8](https://doi.org/10.1007/s00158-018-1912-8). [pp. 21 and 23]
- 85 Justin S. Gray, John T. Hwang, Joaquim R. R. A. Martins, Kenneth T. Moore, and Bret A. Naylor. OpenMDAO: An open-source framework for multidisciplinary design, analysis, and optimization. *Structural and Multidisciplinary Optimization*, 59(4):1075–1104, April 2019. doi:[10.1007/s00158-019-02211-z](https://doi.org/10.1007/s00158-019-02211-z). [p. 23]
- 86 Christopher Lee, David Koo, Karla Telidetzki, Howard Buckley, Hugo Gagnon, and David W. Zingg. Aerodynamic shape optimization of benchmark

- problems using jetstream. In *53rd AIAA Aerospace Sciences Meeting*, 2015. doi:[10.2514/6.2015-0262](https://doi.org/10.2514/6.2015-0262). [p. 25]
- 87 François Bisson and Sivakumaran Nadarajah. Adjoint-based aerodynamic optimization of benchmark problems. In *53rd Aerospace Sciences Meeting*, number AIAA 2015-1948, Kissimmee, Florida, January 2015. doi:[10.2514/6.2015-1948](https://doi.org/10.2514/6.2015-1948).
- 88 Daniel J. Poole, Christian B. Allen, and Thomas C. S. Rendall. Control point-based aerodynamic shape optimization applied to aiaa adodg test cases. In *53rd AIAA Aerospace Sciences Meeting*, 2015. doi:[10.2514/6.2015-1947](https://doi.org/10.2514/6.2015-1947). [p. 25]
- 89 Max M. Munk. The minimum induced drag of aerofoils. Technical Report 121, National Advisory Committee for Aeronautics, 1923. [pp. 25, 27, and 43]
- 90 Stephen C. Smith and Ilan M. Kroo. Computation of induced drag for elliptical and crescent-shaped wings. *Journal of Aircraft*, 30(4):446–452, Jul 1993. doi:[10.2514/3.46365](https://doi.org/10.2514/3.46365). [p. 27]
- 91 Heinrich Blasius. Grenzschichten in flüssigkeiten mit kleiner reibung. *Zeitschrift für angewandte Mathematik und Physik*, 56:1–37, 1908. [p. 29]
- 92 Peter W. Jansen, Ruben E. Perez, and Joaquim R. R. A. Martins. Aerostructural optimization of nonplanar lifting surfaces. *Journal of Aircraft*, 47(5):1490–1503, 2010. [pp. 34 and 74]
- 93 John C. Vassberg, Mark A. DeHaan, S. Melissa Rivers, and Richard A. Wahls. Development of a Common Research Model for applied CFD validation studies. In *26th AIAA Applied Aerodynamics Conference*, 2008. doi:[10.2514/6.2008-6919](https://doi.org/10.2514/6.2008-6919). [p. 48]
- 94 Graeme J. Kennedy and Joaquim R. R. A. Martins. A comparison of metallic and composite aircraft wings using aerostructural design optimization. In *14th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*, Indianapolis, IN, September 2012. doi:[10.2514/6.2012-5475](https://doi.org/10.2514/6.2012-5475). [p. 50]
- 95 Gaetan K. W. Kenway and Joaquim R. R. A. Martins. Buffet-onset constraint formulation for aerodynamic shape optimization. *AIAA Journal*, 55(6):1930–1947, June 2017. doi:[10.2514/1.J055172](https://doi.org/10.2514/1.J055172). [pp. 55, 67, 90, and 99]
- 96 Rhea P. Liem, Gaetan K. W. Kenway, and Joaquim R. R. A. Martins. Multimission aircraft fuel-burn minimization via multipoint aerostructural optimization. *AIAA Journal*, 53(1):104–122, Jan 2015. [pp. 66, 67, and 75]

- 97 Sean Wakayama and Ilan Kroo. Subsonic wing planform design using multidisciplinary optimization. *Journal of Aircraft*, 32(4):746–753, Jul 1995. doi:10.2514/3.46786. [p. 66]
- 98 Howard P. Buckley, Beckett Y. Zhou, and David W. Zingg. Airfoil optimization using practical aerodynamic design requirements. *Journal of Aircraft*, 47(5):1707–1719, Sep 2010. doi:10.2514/1.c000256. [p. 66]
- 99 David J. J. Toal and Andy J. Keane. Efficient multipoint aerodynamic design optimization via cokriging. *Journal of Aircraft*, 48(5):1685–1695, Sep 2011. doi:10.2514/1.c031342. [p. 67]
- 100 Bernard Grossman, Zafer Gürdal, G. J. Strauch, W. M. Eppard, and Raphael T. Haftka. Integrated aerodynamic/structural design of a sailplane wing. *Journal of Aircraft*, 25(9):855–860, 1988. doi:10.2514/3.45670. [p. 74]
- 101 Ian R. Chittick and Joaquim R. R. A. Martins. Aero-structural optimization using adjoint coupled post-optimality sensitivities. *Structures and Multidisciplinary Optimization*, 36:59–70, July 2008. doi:10.1007/s00158-007-0200-9. [p. 74]
- 102 Charles A. Mader, Gaetan K. W. Kenway, Joaquim R. R. A. Martins, and Alejandra Uranga. Aerostructural optimization of the D8 wing with varying cruise Mach numbers. In *18th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*, June 2017. [p. 74]
- 103 Hak tae Lee and Gano B. Chatterji. Closed-form takeoff weight estimation model for air transportation simulation. In *10th AIAA Aviation Technology, Integration, and Operations (ATIO) Conference*, Fort Worth, TX, September 2010. doi:10.2514/6.2010-9156. [p. 79]
- 104 Ilan M. Kroo. *Aircraft Design: Synthesis and Analysis*. Desktop Aeronautics, Palo Alto, CA, September 2006. [p. 79]
- 105 John T. Hwang and Joaquim R. R. A. Martins. A fast-prediction surrogate model for large datasets. *Aerospace Science and Technology*, 75:74–87, April 2018. [p. 83]