

High-Fidelity Aerodynamic Design Optimization of Aircraft Configurations

by

Zhoujie Lu

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Aerospace Engineering)
in The University of Michigan
2014

Doctoral Committee:

Associate Professor Joaquim R.R.A. Martins, Chair
Professor James F. Driscoll
Associate Professor Krzysztof J. Fidkowski
Assistant Professor Kevin J. Maki
J. Douglas McLean, Boeing Commercial Aircraft (Retired)

Research is what I'm doing when I don't know what I'm doing.

—Wernher von Braun

© Zhoujie Lu 2014

All Rights Reserved

To my daughter, Annabelle Lyu, whose birth has brought great motivation and
inspiration to my life.

To my wife, Sandra Liu, whose love and support made this work possible.

ACKNOWLEDGEMENTS

I am fortunate to have had support and guidance from many people during my pursuit of the doctorate degree. First and foremost, I would like to thank my advisor, Professor Joaquim R.R.A. Martins, for his support, encouragement, and guidance throughout my studies at the University of Michigan. His aircraft design class when I was an undergraduate student introduced me to this field. Thank you for always having faith in me and teaching me what it means to be an excellent researcher, instructor, and engineer. I have been truly inspired by your numerous innovative ideas, and your enthusiasm for airplanes throughout these years.

Special thank goes to special committee member Dr. Doug McLean, whose book, *Understanding Aerodynamics: Arguing from the Real Physics* [1], made a significant impact on my work and gave me refreshing insights into aerodynamics and aircraft design. Thanks also go to my dissertation committee, Professor James Driscoll, Professor Krzysztof Fidkowski, Professor Kevin Maki, for their valuable comments and suggestions on my thesis.

I would like to thank my colleagues at the MDOLab for their suggestions and help throughout the years. It has been an absolute pleasure to work with such a bright and friendly group. Special thank goes to Dr. Gaetan Kenway, whose work laid the foundations to this research. This thesis would not have been possible without his guidance, support and expertise. Thank also goes to Yin Yu and Zelu Xu who assisted in plotting the figures, and to David Burdette and John Wurts who helped proofread many of my works.

Special thanks to my friends, classmates and colleagues who have stood with me throughout these years: Dianyun Zhang, Yuntao Chen, Minjie Jiang, Yi Chen, Zhen Wu, Yilong Wen, Chenchen Lu, Hai Wang, Yangbing Lou, Min Jie, Hong Zhou, Yuqiang Bi, Lei Chai, Yue Sun, Clara Ang, Nandan Sawkar, Yin Yu, Zelu Xu, Wenhui Wang, Michael Tian, Yuan Chang, Nansi Xue, John Hwang, Prashanth Rao, Yuxin Luo, Ruben Vandenheede, Gaetan Kenway, Song Chen, Jason Kao, Tristan Dhert, John Wurts, David Burdette, Greame Kennedy, Andrew Lambe, Sandy Mader, Rhea Liem, Edmund Lee, Kentaro Hara, and Kyle Ding. I am grateful for all the wonderful

experiences we have shared.

Finally, I would like to thank my family: my beloved parents, Xiaoping Lu and Xuan Zhou, for always having faith in me and providing me with an exceptional education opportunity, my daughter, Annabelle Lyu, who has made me a morning person and taught me that every day is a surprise, and most importantly, thank goes to my loving wife, Sandra Liu. This thesis would not have been possible without your love, encouragement, and support.

TABLE OF CONTENTS

DEDICATION	ii
ACKNOWLEDGEMENTS	iii
LIST OF FIGURES	ix
LIST OF TABLES	xv
LIST OF APPENDICES	xvii
LIST OF ABBREVIATIONS	xviii
LIST OF SYMBOLS	xx
ABSTRACT	xxii
CHAPTER	
1. Introduction	1
1.1 Motivation	1
1.1.1 Efficient Adjoint Implementation using Automatic Differentiation	1
1.1.2 Aerodynamic Shape Optimization of a Transonic Wing	2
1.1.3 Aerodynamic Design Optimization of a Morphing Trail- ing Edge Wing	5
1.1.4 Aerodynamic Design Optimization of a Blended-Wing- Body Aircraft	6
1.2 Thesis Objectives	8
1.3 Thesis Outline	9
2. Aerodynamic Design Optimization Framework	11
2.1 Geometric Parametrization	11
2.2 Mesh Perturbation	12
2.3 CFD Solver	14

2.4	Optimization Algorithms	15
2.4.1	SNOPT	15
2.4.2	SLSQP	15
2.4.3	PSQP	15
2.4.4	IPOPT	16
2.4.5	CONMIN	16
2.4.6	GCMMA	16
2.4.7	ALPSO	16
2.4.8	NSGA2	16
3.	Adjoint Method Based on	
	Reynolds-Averaged Navier–Stokes Equations	17
3.1	Flow Governing Equations	17
3.2	RANS Automatic Differentiation Adjoint	18
3.2.1	Adjoint Formulation	19
3.2.2	Automatic Differentiation Adjoint	19
3.3	Computation of $\partial\mathcal{R}/\partial w$ and $\partial I/\partial w$	23
3.4	Computation of $\partial\mathcal{R}/\partial x$ and $\partial I/\partial x$	24
3.5	Adjoint Solution Method	25
3.6	Coloring Acceleration Techniques	26
4.	Verifications of Flow Solutions and Derivatives	28
4.1	Verifications of Flow Solutions and Derivatives	28
4.2	Aerodynamic Shape Optimization of an ONERA M6 Wing	30
4.2.1	Verification and Grid Refinement Study	31
4.2.2	Geometric Parametrization, Constraints and Grid Move- ment	34
4.2.3	Optimization Algorithm	36
4.2.4	Computational Resources	36
4.2.5	Optimization Results	38
4.3	Conclusions	47
5.	Gradient-Based and Gradient-Free Optimization	51
5.1	Multi-dimensional Rosenbrock Function	53
5.2	Aerodynamic Twist Optimization	56
5.3	Aerodynamic Shape Optimization	61
5.4	Conclusions	63
6.	Aerodynamic Design Optimization Acceleration Methods	65
6.1	Problem Formulation	66
6.1.1	Baseline Geometry	66

6.1.2	Geometric Parametrization	68
6.1.3	Mesh Convergence Study	68
6.1.4	Optimization Problem Formulation	71
6.1.5	Surface Sensitivity on the Baseline Geometry	73
6.2	Aerodynamic Shape Optimization Results	74
6.2.1	Direct Aerodynamic Shape Optimization	74
6.2.2	Aerodynamic Shape Optimization Using the Richardson's Extrapolation	75
6.2.3	Aerodynamic Shape Optimization Using Multilevel Technique	77
6.3	Conclusions	81
7.	Aerodynamic Design Optimization of a Current-Generation Aircraft	82
7.1	Optimization Problem Formulation	83
7.2	Single-Point Aerodynamic Shape Optimization	83
7.3	Multilevel Optimization Acceleration Technique	87
7.4	Aerodynamic Shape Optimization Starting from a Random Geometry	89
7.5	Effect of the Number of Shape Design Variables	93
7.6	Aerodynamic Shape Optimization without Thickness Reduction	96
7.7	Multipoint Aerodynamic Shape Optimization	97
7.8	Conclusions	101
8.	Aerodynamic Design Optimization of an Adaptive Trailing Edge Wing	104
8.1	Geometric Parametrization	104
8.2	Optimization Problem Formulation	105
8.2.1	Objective Function	105
8.2.2	Design Variables	106
8.2.3	Constraints	106
8.3	Baseline Multipoint Aerodynamic Shape Optimization of the Wing	109
8.4	Morphing Trailing Edge Optimization	110
8.5	Full Wing Morphing Optimization	113
8.6	Comparison between Morphing Trailing Edge and Fully Morphing Wing	114
8.7	Simulating Flights with Morphing Trailing Edge	119
8.8	Conclusions	122
9.	Aerodynamic Design Optimization of a Blended-Wing-Body Aircraft	123

9.1	Geometric Parametrization	124
9.2	Problem Formulation	124
9.3	Initial Geometry	125
9.4	Grid Convergence Study	125
9.5	Optimization Problem Formulation	127
	9.5.1 Objective Function	127
	9.5.2 Design Variables	128
	9.5.3 Constraints	129
9.6	Study 0: Baseline Optimization	132
9.7	Aerodynamic Design Optimization Studies of the BWB	132
	9.7.1 Study 1: Shape and Twist Design Variables	134
	9.7.2 Study 2: Trim Constraint	137
	9.7.3 Study 3: CG Design Variable and Static Margin Constraint	140
	9.7.4 Study 4: Bending Moment Constraint	142
	9.7.5 Study 5: Planform Design Variables	143
	9.7.6 Study 6: Multi-Point Optimization	144
9.8	Conclusions	149
10. Final Remarks		152
	10.1 Conclusions	152
	10.2 Thesis Contributions	155
	10.3 Recommendations	156
APPENDICES		157
BIBLIOGRAPHY		194

LIST OF FIGURES

Figure

1.1	This thesis addresses both methodology and applications of aerodynamic shape optimization.	9
2.1	FFD volume (black) and control surface sub-FFD volume (red) with their respective control points	13
2.2	Sub-FFD volume and control points for a trim control surface deflection of 25 degrees	14
3.1	Euler and RANS flux Jacobian stencil	24
3.2	Euler spatial stencil: 32 cells	25
3.3	Euler state coloring patterns with 13 colors	27
3.4	Euler spatial coloring patterns with 38 colors	27
3.5	RANS state coloring patterns with 35 colors	27
4.1	Volume mesh for bump verification case	29
4.2	C_p distribution of the RANS solution	29
4.3	C_p contours for each grid refinement level compared with experimental data	33
4.4	Polar and grid convergence for each grid level.	35
4.5	FFD control points (blue spheres) and thickness constraints (red lines).	36
4.6	Computational grids used for Euler and RANS analysis. C_p contours are shown for $M = 0.8395$ and $\alpha = 3.06^\circ$	37
4.7	Convergence history of optimality and relative merit function.	39
4.8	Evolution of C_p and cross section shape for section $2z/b = 0.65$	40
4.9	C_p contours for RANS optimized designs.	41
4.10	C_p contours for Euler optimized design.	42
4.11	Cross section shapes for RANS and Euler optimized designs.	43
4.12	C_p contours for baseline and optimized designs for Euler and RANS	48
4.13	Drag divergence curves for three fixed lift coefficients.	49
4.14	Lift distribution for baseline and optimized designs.	50
4.15	Thickness-to-chord ratio and twist distributions for baseline and optimized designs.	50
5.1	Optimization paths for the constrained 2D Rosenbrock function	54
5.2	Gradient-based methods converge faster for the Rosenbrock problem	55
5.3	The gradient-free methods require an excessive number of function evaluations for large numbers of variables.	56

5.4	Visualization of the local and global minimum of 8-dimensional Rosenbrock function	57
5.5	Aerodynamic twist optimization comparison on the L3 grid	58
5.6	Twist optimization convergence history for gradient-free methods for the L3 grid	60
5.7	Twist optimization convergence history for gradient-based methods for the L3 grid	60
5.8	Aerodynamic twist optimization comparison for the L2 grid	61
5.9	Lift distribution comparison of the twist optimized designs	62
5.10	Aerodynamic shape optimization comparison for the L2 grid	63
5.11	Shape optimization convergence history for gradient-based methods for the L2 grid	64
6.1	Baseline CRM wing geometry scaled by its mean aerodynamic chord.	67
6.2	The shape design variables are the z -displacements of 720 FFD control points (red spheres).	68
6.3	The mesh convergence study shows that the difference between the drag values computed with the 28.8 M and the 230.7 M grids is within 1 count.	69
6.4	O-grids of varying sizes were generated using a hyperbolic mesh generator.	70
6.5	Sensitivity study of the baseline wing shows which shape changes yield the largest improvements.	73
6.6	The optimized wing is shock-free and has 8.1% lower drag.	74
6.7	The grid convergence for the shape gradients at different locations on the wing.	77
6.8	The optimization starts with the coarser L2 gradients.	78
6.9	Multilevel optimization with L2 and L1 grids significantly reduced computational cost.	78
6.10	The difference in drag coefficients is within 1 count.	79
6.11	The optimization history of the three approaches.	80
7.1	The optimized wing is shock-free and has 8.5% lower drag.	84
7.2	Most of the computations are performed on the coarse grid.	88
7.3	The optimized results of each grid level exhibit only subtle differences.	88
7.4	The initial geometries are randomly generated from the baseline CRM wing.	90
7.5	The optimization manages to start from a random geometry and converge to an optimal wing that is shock free.	91
7.6	All three optimizations with random starting geometries converged to similar optima.	91
7.7	The merit function values between optimized designs show the local minima.	92
7.8	The Euclidean distances between the multiple local minima are similar and are all under 3.4 in (1.2% of the mean aerodynamic chord).	92
7.9	Optimized designs with varying number of airfoil control points. . .	94
7.10	Optimized designs with varying number of airfoil sections.	95

7.11	Optimized designs for varying numbers of shape design variables. . .	95
7.12	The number of optimization iterations does not decrease significantly as the number of defining airfoils is decreased.	96
7.13	The drag on the optimized wing is five counts higher if no airfoil thickness reduction is allowed.	97
7.14	The multipoint optimization flight conditions represent a five-point stencil in Mach- C_L space.	98
7.15	The multipoint optimized wing has a weak shock on the upper surface for each flight condition.	99
7.16	The multipoint optimized wing has better off-design performance and is more robust to changes in flight conditions.	100
8.1	The wing shape design variables are the z -displacement of 192 FFD control points (red and blue spheres). The trailing edge morphing design variables are only the blue control points.	105
8.2	The multipoint optimized wing has 5.7% lower drag.	110
8.3	Morphing trailing edge optimization at MTOW on-design condition.	111
8.4	Morphing trailing edge optimization at half-weight on-design condition.	112
8.5	Morphing trailing edge optimization at low-Mach low-altitude off-design condition.	112
8.6	Morphing trailing edge optimization at low-Mach high-altitude off-design condition.	113
8.7	Full wing morphing optimization at MTOW on-design condition.	115
8.8	Full wing morphing optimization at low-Mach high-altitude off-design condition.	115
8.9	Drag reduction with morphing TE.	116
8.10	Drag reduction with full morphing wing.	116
8.11	ML/D contour of multipoint optimized baseline wing.	117
8.12	ML/D contour of the morphing TE wing.	117
8.13	ML/D contour of the full morphing wing.	118
8.14	Fuel burn is reduced by 0.7% using morphing TE for DFW–SYD flight.	120
9.1	FFD volume (black) and control surface sub-FFD volume (red) with their respective control points	124
9.2	Sub-FFD volume and control points for a trim control surface deflection of 25 degrees	125
9.3	Geometry of the BWB with the CG location shown in red	126
9.4	Mesh convergence plot of the initial BWB mesh at nominal cruise condition	127
9.5	BWB mesh showing surface and center plane cells	128
9.6	Shape and planform design variables	130
9.7	Study 0: dC_D/dy contour of the baseline BWB	133
9.8	Study 0: dC_L/dy contour of the baseline BWB	133
9.9	Study 1: relative merit function (blue) and optimality (red) history of the optimization	135

9.10	Study 1: optimized pressure distribution, sectional airfoil shape, shock surface, twist, CG, neutral point (NP), and lift distribution of Study 0 (red) and Study 1 (blue) BWB	136
9.11	Study 2: optimized pressure distribution, sectional airfoil shape, shock surface, twist, CG, NP, and lift distribution of Study 1 (red) and Study 2 (blue) BWB	138
9.12	Mach sweep of trimmed and untrimmed designs for Studies 0 and 2	139
9.13	Trim drag of optimized BWB for Studies 0 and 2	139
9.14	Study 3: optimized pressure distribution, sectional airfoil shape, shock surface, twist, CG, NP, and lift distribution of Study 2 (red) and Study 3 (blue) BWB	141
9.15	Study 4: Spanwise lift distribution of optimized designs with various levels of bending moment constraint	143
9.16	Study 5: optimized pressure distribution, sectional airfoil shape, shock surface, twist, CG, NP, and lift distribution of Study 2 (red) and Study 5 (blue) BWB	144
9.17	Study 6: optimized pressure distribution, sectional airfoil shape, shock surface, twist, CG, NP, and lift distribution of Study 1 single-point (red) and Study 6 multi-point (blue) optimized BWB	146
9.18	Study 6: ML/D contours, 99% ML/D , and neutral stability line of twist-optimized baseline BWB	147
9.19	Study 6: ML/D contours, 99% ML/D , and neutral stability line of single-point optimized BWB	147
9.20	Study 6: ML/D contours, 99% ML/D , and neutral stability line of multi-point optimized BWB	148
9.21	Study 6: ML/D contours, 99% ML/D , and neutral stability line of all designs	148
A.1	Wing shape optimization using SNOPT	158
A.2	Wing shape optimization using SLSQP	159
A.3	Wing shape optimization using PSQP	159
A.4	Wing shape optimization using GCMMA	160
B.1	The contour of ML/D	162
B.2	The contour of $\sqrt{ML/D}$	163
B.3	The contour of $aM/c L/D$	163
B.4	Comparison of the maxima locations and 99% contours	164
B.5	$\sqrt{ML/D}$ and $aM/c L/D$ predicts the same maximum peak for $C_L = 0.5$	165
C.1	Morphing trailing edge optimization at $M = 0.85$, $Alt = 28,000$ ft, and $W = 347,500$ kg.	166
C.2	Morphing trailing edge optimization at $M = 0.86$, $Alt = 25,000$ ft, and $W = 347,500$ kg.	167
C.3	Morphing trailing edge optimization at $M = 0.86$, $Alt = 33,000$ ft, and $W = 347,500$ kg.	167
C.4	Morphing trailing edge optimization at $M = 0.85$, $Alt = 34,000$ ft, and $W = 273,200$ kg.	168

C.5	Morphing trailing edge optimization at $M = 0.85$, Alt = 28,000 ft, and $W = 273,200$ kg.	168
C.6	Morphing trailing edge optimization at $M = 0.75$, Alt = 25,000 ft, and $W = 273,200$ kg.	169
C.7	Morphing trailing edge optimization at $M = 0.75$, Alt = 37,000 ft, and $W = 273,200$ kg.	169
C.8	Morphing trailing edge optimization at $M = 0.86$, Alt = 25,000 ft, and $W = 273,200$ kg.	170
C.9	Morphing trailing edge optimization at $M = 0.86$, Alt = 37,000 ft, and $W = 273,200$ kg.	170
C.10	Morphing trailing edge optimization at $M = 0.85$, Alt = 39,000 ft, and $W = 198,900$ kg.	171
C.11	Morphing trailing edge optimization at $M = 0.75$, Alt = 25,000 ft, and $W = 198,900$ kg.	171
C.12	Morphing trailing edge optimization at $M = 0.75$, Alt = 41,000 ft, and $W = 198,900$ kg.	172
C.13	Morphing trailing edge optimization at $M = 0.86$, Alt = 25,000 ft, and $W = 198,900$ kg.	172
C.14	Morphing trailing edge optimization at $M = 0.86$, Alt = 41,000 ft, and $W = 198,900$ kg.	173
C.15	Drag reduction contour for $W = 273,200$ kg with morphing TE. . .	173
C.16	Drag reduction contour for $W = 273,200$ kg with full morphing wing.	174
C.17	Drag reduction contour for $W = 198,900$ kg with morphing TE. . .	174
C.18	Drag reduction contour for $W = 198,900$ kg with full morphing wing.	175
C.19	ML/D contour for $W = 273,200$ kg of multipoint optimized baseline wing.	176
C.20	ML/D contour for $W = 273,200$ kg of the morphing TE wing. . . .	176
C.21	ML/D contour for $W = 273,200$ kg of the full morphing wing. . . .	177
C.22	ML/D contour for $W = 198,900$ kg of multipoint optimized baseline wing.	177
C.23	ML/D contour for $W = 198,900$ kg of the morphing TE wing. . . .	178
C.24	ML/D contour for $W = 198,900$ kg of the full morphing wing. . . .	178
D.1	The BWB mesh(left), FFD volume and 800 shape control points (right)	180
D.2	Case 1: the results of the baseline lift constrained drag minimization with shape variables and geometric constraints. Pressure contour, sectional C_p distribution, airfoil shape, and spanwise lift, t/c , twist distributions of the baseline design (left) and the optimized design (right) are shown.	184
D.3	Case 1: SNOPT optimization convergence history of the feasibility, optimality, and merit function	184

D.4	Case 2: the results of the lift constrained drag minimization with shape variables, geometric, root bending moment, and trim constraints. Pressure contour, sectional C_p distribution, airfoil shape, and spanwise lift, t/c , twist distributions of the baseline design (left) and the optimized design (right) are shown.	186
D.5	Case 2: SNOPT optimization convergence history of the feasibility, optimality, and merit function	187
D.6	Case 3: the results of the lift constrained drag minimization with shape and planform variables, geometric, root bending moment, and trim constraints. Pressure contour, sectional C_p distribution, airfoil shape, and spanwise lift, t/c , twist distributions of the baseline design (left) and the optimized design (right) are shown.	188
D.7	Case 3: SNOPT optimization convergence history of the feasibility, optimality, and merit function	189
D.8	Case 4: the results of the lift constrained drag minimization with shape and planform variables, geometric, root bending moment, trim, and static margin constraints. CG and NP locations, pressure contour, sectional C_p distribution, airfoil shape, and spanwise lift, t/c , twist distributions of the baseline design (left) and the optimized design (right) are shown.	190
D.9	Case 4: SNOPT optimization convergence history of the feasibility, optimality, and merit function	191

LIST OF TABLES

Table

4.1	Accuracy validations of the Euler adjoint	30
4.2	Accuracy validations of the Laminar NS adjoint	30
4.3	Accuracy validations of the frozen-turbulence adjoint	30
4.4	Accuracy validations of the full-turbulence adjoint	31
4.5	ONERA M6 wing mesh sizes	32
4.6	Drag break down for baseline and optimized designs on two mesh levels	44
4.7	Timing breakdown for each optimization	46
4.8	Comparison of aerodynamic coefficients with previous work.	46
5.1	Grid size used in aerodynamic twist optimization	57
5.2	Computational cost comparison of the twist optimization for the L3 grid	59
5.3	Computational cost comparison of the shape optimization for the L2 grid	63
6.1	Mesh convergence study for the baseline CRM wing.	69
6.2	Aerodynamic shape optimization problem	72
6.3	The optimization with the multilevel approach reduces computational time by 84.5%.	80
7.1	The drag differences between the baseline and optimized meshes are nearly constant for each grid level.	86
7.2	The number of iterations on the L0 grid is reduced to 18.	87
7.3	The multiple flight conditions represent a five-point stencil in Mach- C_L space.	98
8.1	The multiple flight conditions represent a five-point stencil in Mach- C_L space and a 2.5 g maneuver case.	106
8.2	Aerodynamic shape optimization problem	108
8.3	The fuel burn reduction is about 1% using morphing TE for various flight trajectories.	121
9.1	Geometric parameters for the BWB	126
9.2	Design variables for the BWB aerodynamic shape optimization . . .	129
9.3	Summary of the constraints used in the BWB aerodynamic shape optimization	131
9.4	Study 1: comparison of twist design variables	137

9.5	Study 3: comparison of optimized aerodynamic coefficients at various CG locations	140
9.6	Study 4: comparison of optimized aerodynamic coefficients at various bending moment constraints	142
9.7	Study 6: Flight conditions for the multi-point optimization	145
9.8	Summary of the results of BWB aerodynamic design optimization studies	149
D.1	Summary of the design variables used in the BWB aerodynamic shape optimization problem	181
D.2	Summary of the constraints used in the BWB aerodynamic shape optimization problem	182
D.3	Case 1: aerodynamic coefficients of the baseline and optimized BWB	183
D.4	Case 2: aerodynamic coefficients of the baseline and optimized BWB	186
D.5	Case 3: aerodynamic coefficients of the baseline and optimized BWB	188
D.6	Case 4: aerodynamic coefficients of the baseline and optimized BWB	190
D.7	Summary of optimal drag coefficients, design variables, and constraints of the aerodynamic shape optimization cases	193

LIST OF APPENDICES

Appendix

A.	Additional Data from Optimizers Comparison Study	158
B.	Comparison of ML/D , \sqrt{ML}/D , and $aM/c L/D$	161
C.	Additional Data from Morphing Trailing Edge Optimization Results .	166
D.	Euler-based Aerodynamic Design Optimization of a Blended-Wing-Body Aircraft	179

LIST OF ABBREVIATIONS

AD	Automatic Differentiation
ADODG	Aerodynamic Design Optimization Discussion Group
ALPSO	Augmented Lagrange multiplier Particle Swarm Optimization
A _o A	Angle of Attack
ASM	Additive-Schwartz Method
BFGS	Broyden-Fletcher-Goldfarb-Shanno algorithm
BM	Bending Moment
BWB	Blended-Wing-Body
CFD	Computational Fluid Dynamics
CG	Center of Gravity
CONMIN	CONstrained function MINimization
CRM	Common Research Model
DDADI	Diagonally Dominant Alternating Direction Implicit method
DOF	Degree of Freedom
DV	Design Variables
FFD	Free Form Deformation
GCMMA	Globally Convergent Method of Moving Asymptotes
GMRES	Generalized Minimum RESidual method
HWB	Hybrid-Wing-Body
IDF	Individual Design Feasible
IPOPT	Interior-Point OPTimizer
ILU	Incomplete LU-factorization
KKT	Karcsch-Kuhn-Tucker
JST	Jameson-Schmidt-Turkel artificial dissipation scheme
LE	Leading Edge
LRC	Long Range Cruise
MDO	Multidisciplinary Design Optimization
MAC	Mean Aerodynamic Chord
MACH	MDO for Aircraft Configurations with High fidelity
MTOW	Maximum Take Off Weight
NK	Coupled Newton-Krylov
NP	Neutral Point
NS	Navier-Stokes
NSGA2	Non Sorting Genetic Algorithm II
OML	Outer Mold Line

ONERA	Office National d'Etudes et Recherches Aérospatiales
PETSc	Portable Extensible Toolkit for Scientific Computing
PSQP	Preconditioned Sequential Quadratic Programming method
PSO	Particle Swarm Optimization
RANS	Reynolds-Averaged Navier Stokes
RK	Runge–Kutta
SA	Spalart–Allmaras
SQP	Sequential Quadratic Programming method
SLSQP	Sequential Least Squares Programming algorithm
SNOPT	Sparse Nonlinear OPTimizer
SUmb	Stanford University multi-block CFD solver
TE	Trailing Edge
TFI	Transfinite Interpolation
TSFC	Thrust Specific Fuel Consumption

LIST OF SYMBOLS

a	Speed of sound
b	Span
C_{bend}	Coefficient of bending moment
C_D	Coefficient of drag
$C_{D_{\text{friction}}}$	Skin friction drag coefficient
$C_{D_{\text{pressure}}}$	Pressure drag coefficient
C_p	Coefficient of pressure
C_M	Coefficient of pitching moment
F_i	Inviscid flux
F_v	Viscous flux
E	Energy
h	Step size
I	Function of interest
K_n	Static margin
M	Mach number
M_∞	Freestream Mach number
ML/D	Mach number times lift over drag ratio
p	Pressure
Pr	Prantl's number
q	Heat flux
\mathcal{R}	Aerodynamic residual
Re	Reynolds number
r	Grid refinement ratio
t	Time
t/c	Thickness-to-chord ratio
u	Velocity
V	Volume
w	Aerodynamic states
x	Design variables
x_{pt}	Nodes on the surface mesh
α	Angle of attack
η_c	Normalized chordwise location
η_s	Normalized spanwise location
τ	Shear stress
μ	Laminar viscosity

μ_t
 Ψ

Turbulence eddy viscosity
Adjoint vector

ABSTRACT

With increasing fidelity and efficiency of numerical simulations, it becomes possible to rely on computational simulations and optimization to achieve a better aircraft design. One of the most computationally intensive disciplines is the aircraft external aerodynamic design. Computational fluid dynamics based on Reynold-averaged Navier–Stokes equations is necessary to accurately resolve the flow field in order to achieve a practical design. High-fidelity CFD poses difficulties to numerical optimization due to its high computational cost, especially when large number of shape design variables are used. This thesis presents an approach to compute the gradients of Reynold-averaged Navier–Stokes equation equations with a Spalart–Allmaras turbulence model using a combination of the adjoint method and automatic differentiation algorithms, for use in gradient-based aerodynamic shape optimization. The resulting gradients are accurate, robust, and efficient. A novel multilevel optimization approach is developed to reduce the computational cost by 84.5%.

With this state-of-the-art Reynolds-averaged Navier–Stokes adjoint and aerodynamic shape optimization framework, we performed three high-fidelity aerodynamic design optimization studies in this thesis. The wing of a Boeing 777-sized aircraft is optimized for single and multiple flight conditions. The drag coefficient is minimized with respect to 720 shape design variables, subject to lift, pitching moment, and geometric constraints, using grids with up to 28.8 M cells. Drag coefficient of the optimized design was reduced by 8.5% relative to the initial design. In addition, the multi-modality of this aerodynamic shape optimization problem is examined by starting optimizations from randomly generated initial geometries. All optimal wings had similar airfoil shapes, with a mean difference of 1.2 in. The variation of the merit function between the multiple local optima confirms that these points are indeed local minima, and indicate that the design space consists of a convex bowl with a small flat bottom that is multimodal. The second application is to optimize the aerodynamics of a near-term aircraft retrofit modification: a wing with morphing trailing edge. A

total of 407 trailing edge optimizations with different Mach numbers, altitudes, and weights, were performed to span the entire cruise flight envelope. A drag reduction in the order of 1% is achieved for on-design conditions, and reductions up to 5% were achieved for off-design conditions. Finally, we extend the aerodynamic shape optimization studies to design an unconventional configuration, the blended-wing-body aircraft. The RANS adjoint and high-fidelity aerodynamic shape optimization framework allowed us to examine the trade-offs between drag coefficient, trim, and static margin of the blended-wing-body configuration. The best compromise between performance and stability was achieved by enforcing a small static margin that can be tolerated in a commercial airplane (1%) and including the center of gravity position as a design variable. This resulted in a trimmed configuration that exhibits a nearly elliptical lift distribution and the lowest drag among the trimmed stable designs. This was achieved by a combination of optimized washout and reflex airfoils.

CHAPTER 1

Introduction

1.1 Motivation

Recent advances in high performance computing have enabled the deployment of full-scale physics-based numerical simulations and optimization in academia and industry. Computational fluid dynamics (CFD) tools and numerical optimization techniques have been widely adopted to shorten the design cycle times and to explore the design space more effectively. High-fidelity methods enable engineers to perform detailed designs earlier in the design process, allowing them to better understand the design trade-offs and make more informed decisions. In addition, advances in sensitivity analysis via the adjoint method [2] has dramatically improved the effectiveness and computational time of aerodynamic shape optimization. However, due to the complexity of the CFD solvers, deployment of adjoint method in Reynolds-averaged Navier–Stokes (RANS) solvers remains a challenging task.

To date, there are only very few examples of solving aerodynamic design problems using a robust and efficient aerodynamic shape optimization algorithm. Therefore, the first goal of this thesis is to develop an efficient aerodynamic shape optimization framework using the adjoint method that enables us to perform high-fidelity large-scale aerodynamic design optimization of aircraft configurations. The second goal is to use this framework to explore a variety of aerodynamic designs, including a current-generation transonic wing, a near-term improvement in the form of a morphing trailing edge wing, and a next-generation blended-wing-body configuration.

1.1.1 Efficient Adjoint Implementation using Automatic Differentiation

There are two types of adjoint approaches: continuous and discrete. In the continuous approach, the adjoint method is directly applied to the governing differential

equations. Partial derivatives of the objectives and residuals with respect to state variables and design variables are combined via the adjoint variables. The governing equations and the adjoint equation along with the boundary conditions are then discretized to obtain numerical solutions. Several authors have demonstrated aerodynamic shape optimization with the continuous adjoint approach [3, 4, 5]. For the discrete adjoint approach, the adjoint method is applied to the set of discretized flow governing equations instead. The gradient produced by the discrete adjoint is exact in the discrete sense and can thus be verified with great precision using the complex-step method [6]. The discrete adjoint approach is also widely used in aerodynamic shape optimization [7, 8, 9, 10, 11, 12, 13, 14]. A comparison between the discrete and continuous adjoint is studied in [15]. Implementation of either continuous or discrete adjoint methods in a complex CFD code remains a challenging and time-consuming task. Derivatives involved in the adjoint formulation for the RANS equations are often difficult to derive and require the manipulation of the governing equations. One way to tackle this difficulty is to use automatic differentiation (AD), either by differentiating the entire code [16, 17], or by selectively differentiating the code to compute the partial derivatives required by the adjoint method [18].

In Chapter 2 to 4 of this thesis, we present an efficient adjoint implementation that extensively uses the automatic differentiation to compute the partial derivative terms for a discrete adjoint of the RANS equations. Simplifications, such as neglecting the turbulent contributions (often called frozen-turbulence), can be made in the formulation. With the automatic differentiation approach, the one-equation Spalart–Allmaras (SA) turbulence model [19] is also differentiated without extensive additional effort. Therefore, an exact linearization to the main flow solution can be achieved. With this implementation, the door to high-fidelity aerodynamic design optimization of aircraft configurations has been opened.

1.1.2 Aerodynamic Shape Optimization of a Transonic Wing

The design of transonic transport aircraft wings is particularly important because of the large number of such aircraft operating on a daily basis, and because small changes in the wing shape may have a large impact on fuel burn. This directly affects both the airlines’ cash operating cost and the emission of green-house gases.

Advances in high-performance computing hardware and algorithms have enabled the ever-increasing fidelity of computational fluid dynamics (CFD) models used for evaluating aircraft performance. As the computational time for a given CFD model reduces, it becomes feasible to use it together with numerical optimization to per-

form aircraft design. While there are various possible optimization techniques, the use of gradient-based algorithms together with an adjoint method that computes the required gradients efficiently has been proven to be particularly effective. Such optimizations typically require a total time equivalent to $\mathcal{O}(10^2)$ CFD simulations to obtain optimal designs. This enables wing designers to shorten design cycle times and thus explore the design space more effectively. They can also obtain detailed designs earlier in the design process, allowing them to better understand the design trade-offs and to make more informed design decisions.

Aerodynamic shape optimization can be dated back to the 16th century, when Newton [20] used calculus of variations to minimize the fluid drag of a body of revolution with respect to the body's shape. Although there were many significant developments in optimization theory after that, it was only in the 1960s that both the theory and the computer hardware became advanced enough to make numerical optimization a viable tool for everyday applications. The application of gradient-based optimization to aerodynamic shape optimization was pioneered in the 1970s. The aerodynamic analysis at the time was a full-potential small perturbation inviscid model, and gradients were computed using finite differences. Hicks *et al.* [21] first tackled airfoil design optimization problems. Hicks and Henne [22] then used a three-dimensional solver to optimize a wing with respect to 11 design variables representing both airfoil shape and the twist distribution.

Because small local changes in wing shape have a large effect on performance, wing design optimization is especially effective with a large number of shape variables. As the number of design variables increases, the cost of computing gradients with finite-differences becomes prohibitive. The development of the adjoint method addressed this issue, enabling the computation of gradients at a cost independent of the number of design variables. For a review of methods for computing aerodynamic shape derivatives, including the adjoint method, see Peter and Dwight [23]. For a generalization of the adjoint method and its connection to other methods for computing derivatives, see Martins and Hwang [24].

Pironneau [25] pioneered the adjoint approach by deriving the adjoint of the Stokes equations and the incompressible Euler equations [26] to optimize airfoil profiles. Jameson [2] extended the adjoint method to handle inviscid compressible flows, making it suitable for the design of transonic airfoils. Since then, adjoint implementations for the compressible Euler equations have been used by various researchers to perform aerodynamic shape optimization. Reuther *et al.* [3, 4], for example, performed the aerodynamic shape optimization of complete aircraft configurations. The devel-

opment of more robust CFD mesh deformations has made it possible to widen the range of the design parameters and study, for example, nonplanar geometries [10, 11]. Hazra *et al.* [27] developed a simultaneous pseudo-timestepping in which stationary states are obtained by solving the non-stationary system of equations representing the state, costate and design equations and successfully applied to optimize an RAE2822 airfoil.

The aerodynamic design of transonic wings requires a model that can represent the shock-wave boundary layer interaction, since there is a strong nonlinear coupling between airfoil shape, wave drag, and viscous effects. Therefore, using a model that relies solely on the Euler equations is insufficient and can even be misleading [28].

Fortunately, the adjoint method has been extended to the compressible Navier–Stokes equations with turbulence models, enabling us to tackle practical aerodynamic design problems. Jameson *et al.* [29] optimized a wing-body configuration modeled with the compressible Navier–Stokes equations using a continuous adjoint approach. They used a 590k-cell mesh and achieved a shock-free solution at Mach 0.86. Anderson and Bonhaus [30] optimized airfoil shapes using a discrete adjoint that included the linearization of the Spalart–Allmaras turbulence model. Nielsen and Anderson [31] further extended the approach to the three-dimensional Reynolds-averaged Navier–Stokes equations. They achieved an 8% drag reduction for the ONERA M6 wing with thickness and camber design variables at two chordwise locations. Brezillon and Dwight [32, 33] optimized the DLR-F6 wing-body configuration using a RANS solver and a discrete adjoint, achieving a 10-count drag reduction by varying 96 design variables. The adjoint approach has also been applied to supersonic flow, such as the sonic boom reduction [34, 35].

The efforts mentioned above use aerodynamic shape optimization frameworks combining different CFD solvers, adjoint implementations, optimizers, and geometry parameterizations, all applied to different design optimization problems. Thus, we need a set of benchmark cases for aerodynamic design optimization, following a similar model to that of the Drag Prediction Workshop [36, 37, 38]. To address this issue, a few researchers formed the Aerodynamic Design Optimization Discussion Group (ADODG) and developed four benchmark problems to test aerodynamic optimization methods. These problems range from the optimization of a two-dimensional airfoil using the Euler equations to three-dimensional shape optimization using the RANS equations.

In Chapter 7, we address the lack of benchmarks in aerodynamic design optimization by presenting a comprehensive set of results for what is currently the most com-

putationally intensive benchmark problem among the test cases: the lift-constrained drag minimization of the NASA Common Research Model (CRM) wing with a RANS model. In addition, we study the effect of the grid size, the number of shape design variables, and their distributions. We also demonstrate the robustness of our aerodynamic shape optimization framework by starting the optimization from a random perturbation of the CRM wing geometry.

1.1.3 Aerodynamic Design Optimization of a Morphing Trailing Edge Wing

Looking beyond the current-generation aircraft wing, one of the fuel-burn reduction strategies that is starting to be used on modern jetliners, such as the Boeing 787, is the use of cruise flaps: A small amount of trailing edge flap and aileron droop is used to optimize the aerodynamic performance at different cruise conditions. The trailing edge devices can alter the spanloads over the course of a flight to reduce drag; however, cruise flaps have limited degrees-of-freedom that may impact the optimal performance. Morphing trailing edge devices, such as those developed by FlexSys, could address this issue. The FlexSys FlexFoil allows independent changes of camber and flap angles at each spanwise location using a smooth morphing surface with no gaps [39, 40]. The morphing trailing edge (TE) has a high technology readiness level and could be retrofitted onto existing aircraft to reduce the drag as much as possible for each flight condition.

Previous studies on morphing trailing edges have focused on the design of the morphing mechanisms, actuators, and structure [39, 40, 41]. In previous aerodynamic studies of the morphing trailing edge, only low-fidelity methods were used [42, 43]. However, small geometry changes, such as the cruise flap extension, require high-fidelity simulations to fully quantify the tradeoff between induced drag and other sources of drag. Therefore, we use a high-fidelity aerodynamic model based on the Reynolds-averaged Navier–Stokes (RANS) equations to examine this tradeoff. The boundary layer is well-resolved and a Spalart–Allmaras turbulence model is used.

Determining the optimal trailing edge shape at each spanwise location for each flight condition is a challenging design task. In this thesis, we use the aerodynamic design optimization framework to explore the optimal design of the morphing trailing edge wing. We use a gradient-based numerical optimization algorithm together with an efficient adjoint implementation [28] to optimize the morphing for the different flight conditions. A database of optimal morphing shapes at different flight conditions is generated using a total of 407 aerodynamic shape optimizations. Once the database

is generated, we can compute the required optimal morphing shapes and related fuel burn reductions for each mission. For comparison purposes, we also perform optimizations with full morphing wings to quantify the theoretical minimal drag for each condition.

1.1.4 Aerodynamic Design Optimization of a Blended-Wing-Body Aircraft

In this thesis, we further extend the aerodynamic optimization studies to a next-generation aircraft. Unconventional aircraft configurations, such as the blended-wing-body (BWB), have the potential to significantly reduce emissions and noise of future large transport aircraft [44]. BWB configuration, also known as the hybrid-wing body (HWB), is characterized by an airfoil-shaped centerbody that integrates payload, propulsion, and control surfaces. Compared to the classic tube-and-wing configuration, the BWB has superior aerodynamic performance [44, 45, 46]: the reduction in the wetted area substantially reduces the skin friction drag, the all-lifting design reduces the wing loading and improves the spanwise lift distribution, the smooth blended wing-centerbody intersection reduces the interference drag, and the area-ruled shape of the BWB reduces the wave drag at high transonic speed. The centerbody provides a substantial portion of the total lift, thus reducing the wing loading. The low wing loading ensures excellent low-speed flight characteristics as well, making heavy high-lift mechanisms, such as double-slotted flaps, redundant. Cross-sectional area of the BWB is similar to that of the Sears–Haack body, which results in lower wave drag at transonic speeds, according to Whitcomb’s area rule [47]. However, the design of BWB configuration introduces new challenges.

The main problem is that since the BWB does not have a horizontal tail, the pressure distributions over the centerbody and wings must be carefully designed to maintain trim and the desired static margin. The thick airfoil shape of the centerbody also makes it a challenge for the BWB to achieve low drag while generating sufficient lift at a reasonable deck angle. Thus, there are critical trade-offs between aerodynamic performance, trim, and stability.

Several authors have investigated the design optimization of the BWB configuration. Liebeck [48, 44] and Wakayama [49, 50] presented the multidisciplinary design optimization (MDO) of the Boeing BWB-450. They used a vortex-lattice model and monocoque beam analysis, and they also considered the trim and stability of the BWB. Qin *et al.* [46, 51] performed an aerodynamic optimization of the European MOB BWB geometry, including inverse design and 3D shape optimization

with a trim constraint. They optimized the design in 3D using Euler-based computational fluid dynamics (CFD). Peigin and Epstein [52] used a genetic algorithm and reduced-order methods to perform a multipoint drag minimization of the BWB with 93 design variables. They used a full Navier–Stokes analysis with reduced-order methods. Kuntawala *et al.* [53, 54] studied BWB planform and shape drag minimization using Euler CFD with an adjoint implementation. Meheut *et al.* [55] performed a shape optimization of the AVECA flying wing planform subject to a low-speed takeoff rotational constraint. They optimized a total of 151 design variables, and they used CFD with a frozen-turbulence (Reynolds-averaged Navier–Stokes) RANS adjoint to compute the gradient. Mader and Martins [56] studied the Euler-based shape optimization of a flying wing considering trim, bending moment constraints, and both static and dynamic stability constraints. Using a minimum induced-drag planform as a reference, they studied the effect of the various constraints on the optimal designs. Their results showed that at subsonic and moderate transonic speeds, the static constraints can be satisfied with airfoil shape variables alone using a reflex airfoil. However, at high transonic speeds, or when considering dynamic stability constraints, the optimal designs required sweep, twist, and airfoil shape variables to minimize the drag while satisfying the constraints. Lyu and Martins [13] investigated the BWB shape optimization with bending moment, trim, and static margin constraints using Euler CFD, including planform optimization. They followed this with a similar study that used a RANS solver [14], which provided the basis for the present study. Reist and Zingg [57] studied the aerodynamic shape optimization of a short-range regional BWB with Euler CFD.

What is missing is a comprehensive and systematic study of a BWB configuration that investigates the design trade-offs between aerodynamic performance, trim, stability, as well as structural considerations, with appropriate fidelity. In this case, the appropriate fidelity is RANS CFD. While Euler-based optimization can provide design insights, the resulting optimal shapes are significantly different from those obtained with RANS, and Euler-optimized shapes tend to exhibit non-physical features, such as a sharp pressure recovery near the trailing edge [28].

Using the optimization framework presented in this thesis, we develop a methodology for the aerodynamic design of BWB configurations that performs optimal trade-offs between the performance and constraints mentioned above, and examines the impact of each constraint on optimal designs. We investigate the design trade-offs by performing a series of aerodynamic shape and planform optimization studies that examine the impact of the design variables and constraints. We explore the effect of

the trim constraint, required static margin, and CG location on the BWB optimal shape. We also investigate the impact of multi-point design optimization.

1.2 Thesis Objectives

The objective of this thesis is to develop a robust and efficient high-fidelity aerodynamic shape optimization framework based on the Reynold-averaged Navier–Stokes equations for solving large-scale aerodynamic design optimization problems. The following tasks are addressed in this thesis.

- Develop an efficient gradient calculation to the RANS equations using the adjoint method and automatic differentiation.
- Verify the aerodynamic and geometry derivatives using the complex-step method.
- Investigate the computational cost of RANS adjoint and frozen-turbulence adjoint.
- Compare the optimum of Euler-based and RANS-based aerodynamic shape optimization.
- Compare the computational cost of aerodynamic shape optimization using gradient-based and gradient-free optimizers.
- Develop an acceleration technique to improve the efficiency of aerodynamic shape optimization.
- Improve the aerodynamic performance of a transonic wing using aerodynamic shape optimization.
- Investigate the local and global optimum in aerodynamic shape optimization.
- Demonstrate the robustness of the aerodynamic shape optimization framework by performing an optimization with randomly generated initial designs.
- Quantify the aerodynamic benefit and obtain the optimal shape of an adaptive trailing edge wing using aerodynamic shape optimization.
- Systematically investigate the aerodynamic design optimization of a blended-wing-body aircraft.
- Investigate the trade-off between aerodynamics and stability of blended-wing-body aircraft.

1.3 Thesis Outline

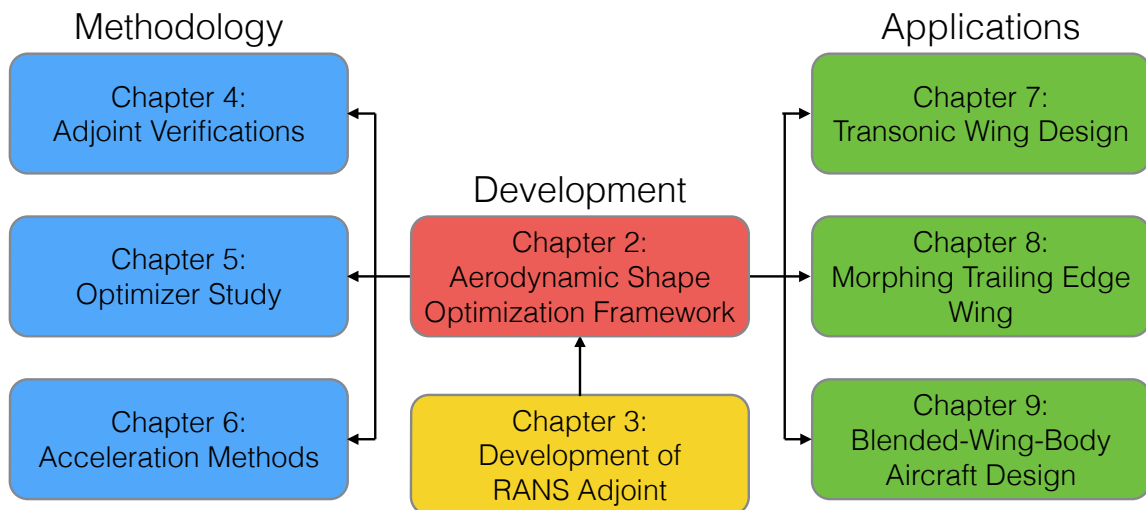


Figure 1.1: This thesis addresses both methodology and applications of aerodynamic shape optimization.

The outline of this thesis is shown in Figure 1.1. This thesis contributes to the state of art of high-fidelity aerodynamic design optimization. Both methodology (shown in blue and yellow) and applications (shown in green) are addressed in the thesis. The aerodynamic shape optimization framework is discussed in Chapter 2. This chapter gives an overview of each component in this framework, including geometry parametrization, mesh perturbation method, and the optimization algorithms. Chapter 3 describes the RANS adjoint implementation using automatic differentiation, and the coloring technique to improve the efficiency of the adjoint. The verification of the aerodynamic and geometric adjoint derivatives using the complex-step is shown in Chapter 4. In this chapter, we also investigate the differences in optimized design and computational cost between Euler-based and RANS-based aerodynamic shape optimization. Once the framework is implemented, we performed a study to compare the gradient-based and gradient-free optimizers. In Chapter 6, we present acceleration techniques to improve the efficiency of aerodynamic design optimization.

The following chapters focus on the applications of high-fidelity aerodynamic design optimization. We begin with applying aerodynamic shape optimization to a transonic wing of a current-generation aircraft in Chapter 7. We use this case to investigate the effect of number of design variables and multi-modality of the problem. We also compare the differences in optimized designs using a single-point and

a multi-point objective function. Once we established this baseline optimization, we further apply the aerodynamic shape optimization to an adaptive morphing trailing edge wing in Chapter 8 to obtain the optimal morphing shape at each flight condition. In addition, we evaluate the drag reduction benefit of the morphing trailing edge wing by comparing that to a non-morphing and a full-morphing wing. Finally, we systematically investigate the aerodynamic design optimization of an unconventional aircraft configuration, the blended-wing-body aircraft in Chapter 9. We examine the trade-offs between aerodynamics, CG location, airfoil and planform shape, and the stability of the aircraft. The conclusions are summarized in Chapter 10.

CHAPTER 2

Aerodynamic Design Optimization Framework

This chapter describes the numerical tools and methods that are used in this thesis. These tools are components of the MDO for Aircraft Configurations with High fidelity (MACH) framework [58, 59, 60], which has been developed over the last decade by the MDOLab at the University of Michigan. MACH can perform the simultaneous optimization of aerodynamic shape and structural sizing variables considering aeroelastic deflections; however, in this thesis we focus solely on the aerodynamic shape optimization.

2.1 Geometric Parametrization

The representation of geometry for aerodynamics requires a complete and watertight description of the Outer Mold Line (OML). The geometries in aerodynamic shape optimization problems are often complex. An exact surface parametrization, such as B-spline surface patches, can result in large number of control points for a complex geometry in order to accurately resolve the geometry. For aerodynamic shape optimization, the initial geometry is usually provided. The optimization only needs to modify the initial geometry to achieve a better objective. There is no need to regenerate the geometry at each iteration. Therefore, we use a geometric parametrization method that only modifies existing geometries, which improves the efficiency of the optimization. A survey of shape parameterization techniques for high-fidelity multidisciplinary shape optimization can be found in [61].

We choose to use a free form deformation (FFD) approach to parametrize the geometry. The geometric parametrization in this framework was developed by Kenway [62]. The FFD volume parametrizes the geometry changes rather than the geometry itself, resulting in a more efficient and compact set of geometry design variables, thus making it easier to handle complex geometric manipulations. Any geometry

may be embedded inside the volume by performing a Newton search to map the parameter space to physical space. All the geometric changes are performed on the outer boundary of the FFD volume. Any modification of this outer boundary can be used to indirectly modify the embedded objects. The key assumption of the FFD approach is that the geometry has constant topology throughout the optimization process, which is usually the case for wing design. In addition, since FFD volumes are tri-variate B-spline volumes, the sensitivity information of any point inside the volume can be easily computed. By moving a set of control points together, large geometric variations such as span and sweep can be produced. The control points can also be moved individually to produce airfoil shape changes. Another less technical way to understand FFD approach is to imagine the FFD volume as a block of jello. If you push the outside of the jello, the inside of the jello will change its shape as well. Figure 2.1 shows the FFD volume and geometric control points for a blended-wing-body aircraft. This optimization formulation will be discussed in more detail in Chapter 9.

The FFD approach can be extended to include nested FFD volumes. Those sub-FFD can be used to provide additional degrees of freedom to perform local shape changes. In the BWB example, we use nested FFD to simulate control surfaces movement on the rear centerbody, which are analogous to elevators on a conventional configuration. The result is a sub-FFD that is embedded in the main FFD. Any changes in the main FFD are propagated to the sub-FFD. The sub-FFD is set to rotate about the hinge line of the control surface. When the sub-FFD rotates, the embedded geometry changes the local shape accordingly. Because of the constant topology assumption of the FFD approach, and the limitation of the mesh perturbation, the surface has to be continuous around the control surfaces, eliminating the elevator gap. Therefore, when the control surfaces deflect, there is a transition region between the control surface and the centerbody, similar to those studied in a continuous morphing wing [40]. Figure 2.2 shows the sub-FFD volume and the geometry, with a trim control surface deflection of 25 degrees.

2.2 Mesh Perturbation

Since FFD volumes modify the geometry during the optimization, we must perturb the mesh for the CFD analysis to solve for the modified geometry. The mesh perturbation scheme used in this work is a hybridization of algebraic and linear elasticity methods developed by Kenway [62]. The idea behind the hybrid warping scheme is to




Figure 2.1: FFD volume (black) and control surface sub-FFD volume (red) with their respective control points

apply a linear-elasticity-based warping scheme to a coarse approximation of the mesh to account for large, low-frequency perturbations, and to use the algebraic warping approach to attenuate small, high-frequency perturbations. The goal is to compute a high-quality perturbed mesh similar to that obtained using a linear elasticity scheme but at a much lower computational cost.

The algebraic mesh perturbation used here is an algebraic transfinite interpolation (TFI) [63]. Linear blending functions and linear interpolation are used. The linear elasticity method is based on a linear-spring analogy [64]. Each edge of the grid is considered as a linear spring with a stiffness inversely proportional to its length. A linear system is formed to describe the displacement of the coarsened mesh nodes. The displacement at the surface is prescribed from the design variables. The mesh nodes at the symmetry plane are constrained to remain on that plane. Finally, the nodes at the far field are constrained to be fixed. This hybrid mesh perturbation method is robust and efficient for the applications in this thesis. More detail of the




Figure 2.2: Sub-FFD volume and control points for a trim control surface deflection of 25 degrees

mesh perturbation method can be found in [62].

2.3 CFD Solver

We use the Stanford University multiblock (SUmB) [65] flow solver. SUmB is a finite-volume, cell-centered multiblock solver for the compressible Euler, laminar Navier–Stokes, and RANS equations (steady, unsteady, and time-periodic). It provides options for a variety of turbulence models with one, two, or four equations and options for adaptive wall functions. In this thesis, we focus on solving the RANS equation with the Spalart–Allmaras turbulence model. The Jameson–Schmidt–Turkel (JST) scheme [66] augmented with artificial dissipation is used for the spatial discretization. The main flow is solved using an explicit multi-stage Runge–Kutta method along with a geometric multi-grid scheme. A segregated Spalart–Allmaras (SA) turbulence equation is iterated with the diagonally dominant alternating direction implicit (DDADI) method. An automatic differentiation adjoint for the Euler and RANS equations was developed to compute the gradients [18, 28]. The adjoint implementation supports both the full-turbulence and frozen-turbulence modes. The adjoint equations are solved with preconditioned GMRES [67] using PETSc [68, 69, 70]. More detail of the adjoint is discussed in Chapter 3.

2.4 Optimization Algorithms

Because of the high computational cost of CFD solutions, it is critical to choose an efficient optimization algorithm that requires a reasonably low number of function calls. Gradient-free methods, such as genetic algorithms, have a higher probability of getting close to the global minimum for cases with multiple local minima. However, slow convergence and the large number of function calls make gradient-free aerodynamic shape optimization infeasible with the current computational resources, especially for large numbers of design variables. Therefore, for most of the optimization in this thesis, we use a gradient-based optimizer combined with adjoint gradient evaluations to solve the problem efficiently. More detail of the optimizer studies are discussed in Chapter 5.

2.4.1 SNOPT

We use SNOPT (sparse nonlinear optimizer) [71] through the Python interface pyOpt [72] for most of the optimization in this thesis. SNOPT is a gradient-based optimizer that implements a sequential quadratic programming method; it is capable of solving large-scale nonlinear optimization problems with thousands of constraints and design variables. SNOPT uses a smooth augmented Lagrangian merit function, and the Hessian of the Lagrangian is approximated using a limited-memory quasi-Newton method.

For comparison purpose, a number of other gradient-based and gradient-free optimizers are also used in this thesis. Those optimizers are described below.

2.4.2 SLSQP

SLSQP is a sequential least squares programming algorithm [73] that evolved from the least squares solver of Lawson and Hanson [74]. The optimizer uses a quasi-Newton Hessian approximation and an L1-test function in the line search algorithm. Kraft [75] also applied this method to aerodynamic and robotic trajectory optimization.

2.4.3 PSQP

PSQP is a preconditioned SQP method with a BFGS variable metric update. It can handle large scale problems with nonlinear constraints.

2.4.4 IPOPT

IPOPT implements a primal-dual interior-point algorithm with a filter line search method [76]. The barrier problem is solved using a damped Newton's method. The line search method includes a second order correction.

2.4.5 CONMIN

CONMIN solves linear or nonlinear optimization problems using the method of feasible directions [77]. It minimizes the objective function until it reaches an infeasible region. The optimization then continues by following the constraint boundaries in a descent direction.

2.4.6 GCMMA

GCMMA is a modified version of the method of moving asymptotes, designed for nonlinear programming and structural optimization [78]. It solves a strictly convex approximating sub-problem at each iteration. GCMMA guarantees convergence to a local minimum from any feasible starting point.

2.4.7 ALPSO

ALPSO is a parallel augmented Lagrange multiplier particle swarm optimization (PSO) solver written in Python [79]. This method takes advantage of PSO methods, which can solve non-smooth objective functions and is more likely to find the global minimum. Augmented Lagrange multipliers are used to handle constraints. ALPSO can be used for nonlinear, non-differentiable, and non-convex problems. Perez and Behdinan [80] applied this method to a non-convex, constrained structural problem. Other applications include aerostructural optimization of nonplanar lifting surfaces [81] and aeroservoelastic design optimization of a flexible wing [82].

2.4.8 NSGA2

NSGA2 is a non-dominant sorting based multi-objective evolutionary algorithm [83]. The optimizer enforces constraints by tournament selection. It can solve non-smooth and non-convex multi-objective functions and tends to approach the global minimum.

CHAPTER 3

Adjoint Method Based on Reynolds-Averaged Navier–Stokes Equations

This chapter presents an approach for the rapid implementation of an adjoint solver for the Reynolds-averaged Navier–Stokes equations with a Spalart–Allmaras turbulence model. This work is based on a previous paper presented by the author [28]. Automatic differentiation is used to construct the partial derivatives required in the adjoint formulation. The resulting adjoint implementation is computationally efficient and highly accurate. The assembly of each partial derivative in the adjoint formulation is discussed. In addition, a coloring acceleration technique is presented to improve the adjoint efficiency.

To develop the RANS adjoint, it is necessary to have a thorough understanding of the governing equations and the flow solvers, such as the size of stencils, the vector of state variables, the call sequences, etc. In this section, we discuss the backgrounds, methods, and tools that are involved in the formulation and implementation of the RANS adjoint.

3.1 Flow Governing Equations

The RANS equations are a set of conservation laws that relate mass, momentum, and energy in a control volume. To simplify the expressions, the RANS equations, (3.1) are demonstrated in 2 dimensions.

$$\frac{\partial w}{\partial t} + \frac{1}{A} \oint F_i \cdot \hat{n} dl - \frac{1}{A} \oint F_v \cdot \hat{n} dl = 0 \quad (3.1)$$

The state variable w , inviscid flux F_i and viscous flux F_v are defined as follows.

$$w = \begin{bmatrix} \rho \\ \rho u_1 \\ \rho u_2 \\ \rho E \end{bmatrix} \quad (3.2) \quad F_{i_1} = \begin{bmatrix} \rho u_1 \\ \rho u_1^2 + p \\ \rho u_1 u_2 \\ (E + p)u_1 \end{bmatrix} \quad (3.3) \quad F_{v_1} = \begin{bmatrix} 0 \\ \tau_{11} \\ \tau_{12} \\ u_1 \tau_{11} + u_2 \tau_{12} - q_1 \end{bmatrix} \quad (3.4)$$

The shear stress and heat flux depends on both the laminar viscosity μ and the turbulent eddy viscosity μ_t , where

$$\tau_{11} = (\mu + \mu_t) \frac{M_\infty}{Re} \frac{2}{3} (2u_1 - u_2), \quad (3.5)$$

$$q_1 = -\frac{M_\infty}{Re(\gamma - 1)} \left(\frac{\mu}{Pr} + \frac{\mu_t}{Pr_t} \right) \frac{\partial a^2}{\partial x_1}. \quad (3.6)$$

The laminar viscosity is determined by Sutherland's law. The turbulent eddy viscosity can be updated with turbulence models. In this thesis, we use SA turbulence model and it is solved in a segregated fashion to update the turbulent eddy viscosity. When solving for the main flow variables at each iteration, the turbulence variables are frozen, and vice versa.

3.2 RANS Automatic Differentiation Adjoint

High fidelity aerodynamic shape optimization with large number of design variables requires an efficient gradient calculation. Traditional methods, such as finite-difference, are straightforward to implement, but are inefficient for large-scale optimization problems and are subject to subtractive cancellation error. The complex-step method alleviates the errors resulting from subtractive cancellation and can provide machine-accurate gradients, but similar to finite differencing, is inefficient for large number of design variables. The total number of function evaluation scales with the number of the design variables. Thus, for optimization problems with large number of design variables, the cost of one complete derivative computation can be on the order of hundreds or thousands flow solutions, which is generally prohibitive when using high-fidelity models. For the adjoint method, the cost of the derivative computations is nearly independent of the number of design variables and scales only with the number of functions of interest, which is generally much smaller than the number of design variables.

3.2.1 Adjoint Formulation

For a CFD solver, the discrete adjoint equations can be expressed as,

$$\left[\frac{\partial \mathcal{R}}{\partial w}\right]^T \Psi = -\frac{\partial I}{\partial w}, \quad (3.7)$$

where \mathcal{R} is the residual of the computation, w is the state variables, I is the function of interest, and Ψ is the adjoint vector. We can see that the adjoint equations do not involve the design variable x . For each function of interest, the adjoint vector only needs to be computed once, and it is valid for all design variables. Once the adjoint vector is computed, the total derivatives can be obtained using the total derivative equation,

$$\frac{dI}{dx} = \frac{\partial I}{\partial x} - \frac{\partial I}{\partial w} \left[\frac{\partial \mathcal{R}}{\partial w}\right]^{-1} \frac{\partial \mathcal{R}}{\partial x} = \frac{\partial I}{\partial x} + \Psi^T \frac{\partial \mathcal{R}}{\partial x}. \quad (3.8)$$

Partial derivatives in the equations represent an explicit dependence that do not require convergence of the residual. In the case of CFD, by using the adjoint method, the cost of the total derivatives of any number of design variables can be on the order of or less than the cost of one single CFD solution. Large-scale engineering optimization problems can be solved within a reasonable time only with efficient gradient calculation.

3.2.2 Automatic Differentiation Adjoint

With the adjoint formulation, there is still one challenge — computing the partial derivatives efficiently. One can naïvely use finite difference or complex-step to compute these partial derivatives. However, the prohibitive computational cost from using those methods would completely defeat the purpose of the adjoint method. The partial derivatives can also be derived analytically by hand, but such derivation is non-trivial for a complex CFD solver and typically requires a lengthy development time for the adjoint.

In order to counter those disadvantages, Mader and Martins [18] proposed the ADjoint approach. The main idea is to utilize automatic differentiation techniques to compute partial derivative terms for the adjoint method. The automatic differentiation approach, also known as computational differentiation or algorithmic differentiation, relies on a tool to perform source code transformation on the original solver to create the capability of computing derivatives. This method is based on a systematic application of the chain rule to each line of the source code. There are two

modes in automatic differentiation: forward mode and reverse mode. The forward mode propagates the derivatives along the execution path of the original code. The reverse mode first follows that execution path and in the *store-all* approach stores all the intermediate variables. The source code is then re-run in the *reverse* execution order and the stored intermediate variables are used in the linearization of each line of code.

We use the following example to demonstrate the underlining methodology of forward and reverse mode automatic differentiation,

$$f(x_1, x_2) = x_1x_2 + \sin(x_1). \quad (3.9)$$

This function can be written as the sequence of elementary operations on the intermediate variables q_i resulting in the following sequence,

$$\begin{aligned} q_1 &= x_1 \\ q_2 &= x_2 \\ q_3 &= q_1q_2 \\ q_4 &= \sin(q_1) \\ q_5 &= q_3 + q_4 \end{aligned} \quad (3.10)$$

This sequence is then used to compute the derivative of Equation (3.9).

Forward Mode Forward mode AD is simpler and more intuitive of the two approaches. In Equation (3.9), assuming that x_1 and x_2 are independent inputs, the rules of differentiation are applied to the sequence in Equation (3.10) as follows.

$$\begin{aligned} \Delta q_1 &= \Delta x_1 \\ \Delta q_2 &= \Delta x_2 \\ \Delta q_3 &= \Delta q_1q_2 + q_1\Delta q_2 \\ \Delta q_4 &= \cos(q_1)\Delta q_1 \\ \Delta q_5 &= \Delta q_3 + \Delta q_4 \end{aligned} \quad (3.11)$$

Once the sequence and its corresponding gradients for the function in Equation (3.9) are known, x_1 and x_2 can be seeded to determine the gradient of the function. Since x_1 and x_2 are assumed to be independent inputs, seeding each input independently means to set the variation to one while the other remains zero such that $\Delta x_1 = [1, 0]$

and $\Delta x_2 = [0, 1]$. Forward mode AD sweeps over the computations in Equation (3.11) twice, once for each input, and adds the separate derivative evaluations as follows.

$$\begin{aligned}
\Delta f(x_1, x_2) &= [f_{x_1}, f_{x_2}] \\
&= \Delta q_{5_{x_1}} + \Delta q_{5_{x_2}} \\
&= (\Delta q_{3_{x_1}} + \Delta q_{4_{x_1}}) + (\Delta q_{3_{x_2}} + \Delta q_{4_{x_2}}) \\
&= ((1)q_2 + q_1(0) + \cos(q_1)(1)) + ((0)q_2 + q_1(1) + \cos(q_1)(0)) \\
&= q_2 + \cos(q_1) + q_1 \\
&= x_1 + x_2 + \cos(x_1)
\end{aligned} \tag{3.12}$$

This is the expected derivative for the original function in Equation (3.9), and can be written in a more general format by considering the general sequence $q = (q_1, \dots, q_n)$. Considering m input variables and p output variables, the sequence becomes $q = (q_1, \dots, q_m, q_{m-p+1}, \dots, q_n)$. For $i > m$, each q_i must have a dependence on some member of the sequence prior to i . If $k < i$, the entry q_i of the sequence must depend explicitly on q_k . The forward mode can then be written as the chain rule summation as shown in [84],

$$\Delta q_i = \sum \frac{\partial q_i}{\partial q_k} \Delta q_k, \tag{3.13}$$

where $i = m+1, \dots, n$ and $k < i$. The forward mode AD evaluates the gradients of the intermediate variables first such that $\Delta q_1, \dots, \Delta q_{i-1}$ are known prior to the evaluation of Δq_i . It is easy to see that the forward mode builds up the derivative information as it progresses forward through the algorithm, producing the derivative information for all of output variables with respect to a single seeded input variable.

Reverse Mode The reverse mode, though less intuitive, is dependent only on the number of outputs. With reference to the previous example, it is easier to understand reverse mode AD by examining the partial derivatives of f . Consider the following,

$$\frac{\partial q_5}{\partial q_1} = \frac{\partial q_5}{\partial q_1} \frac{\partial q_1}{\partial q_1} + \frac{\partial q_5}{\partial q_2} \frac{\partial q_2}{\partial q_1} + \frac{\partial q_5}{\partial q_3} \frac{\partial q_3}{\partial q_1} + \frac{\partial q_5}{\partial q_4} \frac{\partial q_4}{\partial q_1}, \tag{3.14}$$

where q_5 represents the single output f . The reverse mode first runs a forward sweep to determine all of the intermediate values in the sequence. Then, starting with a single output variable, (q_5 in this case), the AD tool steps *backwards* through the algorithm to compute the derivatives in reverse order. The example, and the implementation of the sequence in Equation (3.14), produces the following final result.

$$\begin{aligned}
\frac{\partial q_5}{\partial q_5} &= 1 \\
\frac{\partial q_5}{\partial q_4} &= 1 \\
\frac{\partial q_5}{\partial q_3} &= 1 \\
\frac{\partial q_5}{\partial q_2} &= \frac{\partial q_5}{\partial q_3} \frac{\partial q_3}{\partial q_2} = (1)(q_1) \\
\frac{\partial q_5}{\partial q_1} &= \frac{\partial q_5}{\partial q_3} \frac{\partial q_3}{\partial q_1} + \frac{\partial q_5}{\partial q_4} \frac{\partial q_4}{\partial q_1} = (1)(q_2) + (1)(\cos(q_1))
\end{aligned} \tag{3.15}$$

where we have,

$$\begin{aligned}
\frac{\partial q_5}{\partial q_1} &= \frac{\partial f}{\partial x_1} = x_2 + \cos(x_1) \\
\frac{\partial q_5}{\partial q_2} &= \frac{\partial f}{\partial x_2} = x_1.
\end{aligned} \tag{3.16}$$

The advantage here is that only a single reverse sweep is required to evaluate the derivatives with respect to both x_1 and x_2 . Should there be a greater number of inputs, which is typical in an aerodynamic shape optimization problem, a single forward sweep to accumulate the code list as well as a single reverse mode sweep is all that would be necessary to calculate the sensitivities for a single output.

The disadvantage of the reverse mode is that the implementation is more complicated than the forward mode. The reverse mode was used in the original development of the ADjoint method and so is used as a benchmarking tool in the development of the forward mode Adjoint method. To avoid the high computational costs associated with using the forward mode of AD, a coloring method is used to accelerate the computation.

For the adjoint equations (3.7) and (3.8), the partial derivatives $\partial \mathcal{R} / \partial w$, $\partial I / \partial w$, $\partial I / \partial x$, and $\partial \mathcal{R} / \partial x$ are computed with forward automatic differentiation. The following sections explain the implementation of each of the partial derivatives in detail. There are various automatic differentiation tools available including ADIC [85], AD-IFOR [86], FADBAD++ [87], OpenAD/F [88], and TAPENADE [89]. The work presented in this thesis uses TAPENADE to perform the task. TAPENADE is an automatic differentiation engine developed by the Tropics team at Institut National de Recherche en Informatique et Automatique and supports both forward and reverse modes.

3.3 Computation of $\partial\mathcal{R}/\partial w$ and $\partial I/\partial w$

To compute $\partial\mathcal{R}/\partial w$, there are three flux calculations involved: inviscid fluxes, artificial dissipation fluxes, and viscous fluxes. For Euler equations, the combined stencil is the current cell and the 12 adjacent cells in each of the three dimensions — a total of 13 cells, as shown in Figure 3.1(a). The laminar and RANS equations have much larger stencils due to the nodal averaging procedure used in the viscous fluxes. The RANS stencil is a dense 3x3x3 block around the center cell plus additional 6 adjacent cells in each direction, as shown in Figure 3.1(b). The size and the shape of the stencil is important for the coloring acceleration techniques, which is discussed further in Section 3.6. For RANS equations, the state vector w contains the five main flow state variables and one turbulence variable for the one-equation SA model. Therefore, the residual computation for the SA equation also needs to be included in the automatic differentiation. The contribution of the turbulence to the main flow residual is included via the turbulence variable. The frozen-turbulence assumption can be made by neglecting the turbulence contribution to the main flow. Since $\partial\mathcal{R}/\partial w$ is a square matrix, in principle both forward and reverse modes would require similar number of function calls to form the matrix. However, forward mode is more intuitive and has lower overhead cost and for forming $\partial\mathcal{R}/\partial w$, the forward mode is faster than the reverse mode in practice. $\partial\mathcal{R}/\partial w$ is stored in transpose form in a block compressed sparse row matrix format.

Special care must be taken for the computation of $\partial I/\partial w$ with forward mode AD. If the routine to compute I , which we will assume consists of integrated forces or moments on wall boundary, is simply included with the residual evaluation, all cells near the surface that influence the force evaluation on the wall would have to be perturbed independently and the advantage of the graph coloring approach described in Section 3.6 would be nullified.

To enable the evaluation of $\partial I/\partial w$ simultaneously with $\partial\mathcal{R}/\partial w$ it is necessary to evaluate individual forces and moments at each cell, not just the overall sum. Stencils for individual force and moment computations are compact. For both Euler and RANS cases this force stencil is smaller than the corresponding residual stencil. For the linear pressure extrapolation wall boundary condition, the Euler force stencil has only two cells: the cell on the surface and the cell above. The RANS force stencil consists of a 3x3 patch on the surface and one layer above, with a total of 18 cells. Both Euler and RANS force stencil can be packed inside the respective flux Jacobian stencils. Once the individual forces are resolved, their contribution to the chosen

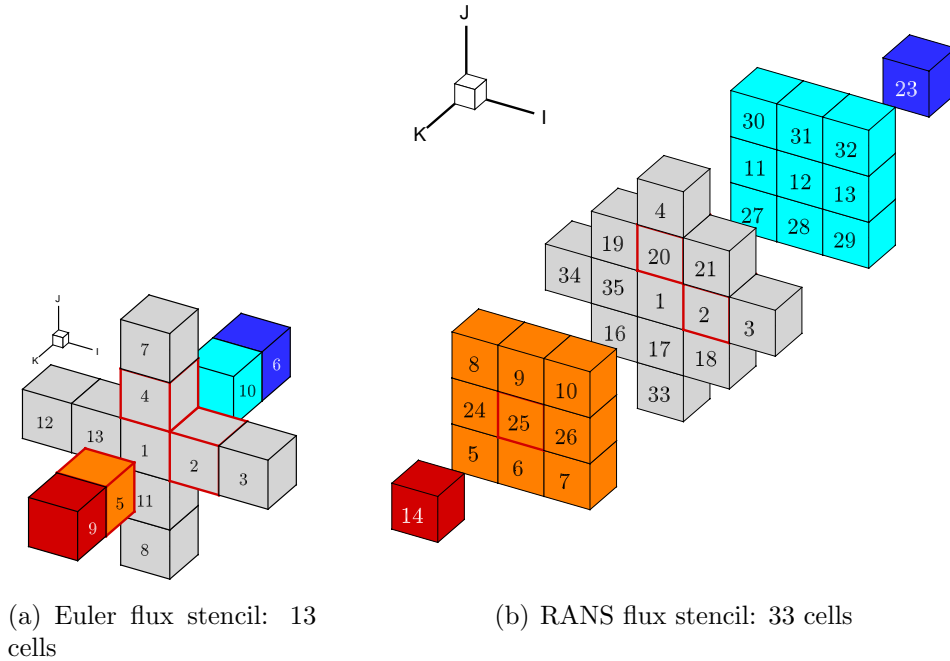


Figure 3.1: Euler and RANS flux Jacobian stencil

objective, I , is evaluated and the correct contribution can be added to $\partial I/\partial w$.

3.4 Computation of $\partial \mathcal{R}/\partial x$ and $\partial I/\partial x$

The calculations of $\partial \mathcal{R}/\partial x$ and $\partial I/\partial x$ depend on the design variables. For aerodynamic shape optimization, we are generally interested in geometric design variables, such as airfoil profile, wing twists, etc, and flow condition design variables, such as Mach number, angle of attack, side-slip angle etc. The partials that involve flow design variables are relatively straight-forward. Each flow design variables are seeded and forward mode AD is used to obtain the residual and objective partial derivatives. No coloring scheme is necessary, since only one pass of the residual routine is needed for each flow design variables.

The partial derivatives for the geometric design variables require careful implementation. In an effort to modularize codes, Sumb does not require the specific information about the geometric design variables. Instead, we calculate $\partial \mathcal{R}/\partial x_{pt}$ and $\partial I/\partial x_{pt}$, where x_{pt} includes all the nodes in the CFD mesh. We use a separate utility to perform the mesh deformation sensitivity calculation [62] and manipulation of the surface geometry. The surface geometry is manipulated using the free-form deformation (FFD) approach. The spatial derivatives are then transformed into the

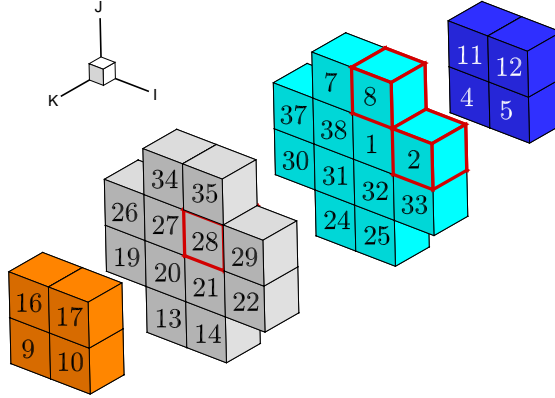


Figure 3.2: Euler spatial stencil: 32 cells

derivatives with respect to the control points of the FFD volume.

To compute $\partial\mathcal{R}/\partial x_{pt}$ and $\partial I/\partial x_{pt}$, we again use forward mode AD. The choice of forward mode may not be obvious here. The benefit of using forward mode AD is that the same Sumb residual routine can be used in both state and spatial derivatives, resulting a much less demanding implementation and fewer modifications. Similar to the state partial derivatives, both $\partial\mathcal{R}/\partial x_{pt}$ and $\partial I/\partial x_{pt}$ can be computed at the same time. The center of the stencils for the spatial derivatives is a node instead a cell. Figure 3.2 shows the Euler $\partial\mathcal{R}/\partial x_{pt}$ stencil, with a total of 32 cells. The corresponding RANS stencil is a dense $4\times 4\times 4$ block containing 64 cells. The Euler stencil for $\partial I/\partial x_{pt}$ is a 4×4 surface patch, while the RANS one includes one addition layer above. As we can see, both spatial force derivatives can also be fitted inside the spatial residual stencils.

3.5 Adjoint Solution Method

We use a preconditioned GMRES [67] algorithm in PETSc (Portable, Extensible Toolkit for Scientific Computation) [68, 69, 70] to solve the adjoint system. PETSc is a suite of data structures and routines for the scalable parallel solution of scientific applications modeled by partial differential equations. The system is preconditioned with restrictive additive Schwarz method and incomplete LU (ILU) factorization on each sub-domain. We found that GMRES with approximate preconditioner produced using a first-order approximation is very effective with Euler adjoint. The RANS adjoint, especially without the frozen-turbulence assumption, is considerably more stiff than the Euler adjoint with the problem more prominent at high Reynolds number. A stronger preconditioner, such as full $\partial\mathcal{R}/\partial w$ Jacobian as preconditioner, may be

necessary.

3.6 Coloring Acceleration Techniques

As previously noted, a naïve approach for computing $\partial\mathcal{R}/\partial w$ and $\partial\mathcal{R}/\partial x$ would require a total of $N_{\text{state}} + 3 \times N_{\text{nodes}}$ evaluations, where N_{state} and N_{nodes} are the total number of cells and nodes respectively. In this approach each column (or row of the transposed Jacobian) is computed one at a time. If we however, exploit the sparsity structure of $\partial\mathcal{R}/\partial w$ and $\partial\mathcal{R}/\partial x$, it is possible to fully populate the Jacobians with far fewer function evaluations. The general idea is to determine groups of independent columns of the Jacobian. A group of columns is considered independent if no row contains more than one nonzero entry. This allows a group of independent columns to be evaluated simultaneously. The process of determining which columns are independent is known as graph coloring. The determination of an optimal (smallest) set of *colors* for a general graph is quite challenging. For unstructured grids, a greedy coloring scheme can be used resulting in a satisfactory number of colors[90].

For structured grids with regular repeating stencils, the graph coloring problem is substantially simpler [91]. Consider the 13-cell stencil for the Euler residual evaluation shown in Figure 3.1(a). It is clear at least 13 colors will be required. Determining the optimum graph coloring for this case is equivalent to finding a three dimensional packing sequence that minimizes the unused space between stencils. Fortunately, for this stencil, a perfect packing sequence is possible and precisely 13 colors are required. A three-dimensional view of the stencil packing is shown in Figure 3.3. For the 33 cell RANS stencil, it is not possible to perfectly pack the stencils. We have, however, found a coloring scheme that requires only 35 colors, shown in Figure 3.4. To assign the coloring number mathematically to each cell, simple formula can be derived using the remainder function $\text{mod}(m, n)$. m is a function determined by the numbered stencil and n is the total number of colors required to populate the matrix. Equations (3.17) through (3.19) show the coloring function for the Euler and RANS stencils. These optimal graph colorings reduce the number of forward mode AD perturbations to a fixed constant, independent of the mesh size.

$$C_{\text{Euler, state}}(i, j, k) = \text{mod}(i + 3j + 4k, 13) \quad (3.17)$$

$$C_{\text{Euler, spatial}}(i, j, k) = \text{mod}(i + 7j + 27k, 38) \quad (3.18)$$

$$C_{\text{RANS, state}}(i, j, k) = \text{mod}(i + 19j + 11k, 35) \quad (3.19)$$




Figure 3.3: Euler state coloring patterns with 13 colors




Figure 3.4: Euler spatial coloring patterns with 38 colors




Figure 3.5: RANS state coloring patterns with 35 colors

CHAPTER 4

Verifications of Flow Solutions and Derivatives

In this chapter, the RANS adjoint discussed previously is verified with complex-step method using a flow over a bump case. We further verify the derivatives and demonstrate the aerodynamic shape optimization capability by performing a standard optimization test case: aerodynamic shape optimization of an ONERA M6 wing. The results are compared with Euler-based aerodynamic shape optimization and previous work. Finally, the effects of the frozen-turbulence assumption on the accuracy and computational cost are assessed. This work is based on a previous paper presented by the author. [28]

4.1 Verifications of Flow Solutions and Derivatives

A flow over a bump case is chosen as the test case to verify Euler, Laminar NS, and RANS adjoint solutions. The computational mesh for the test case used is shown in Figure 4.1. It is a single block mesh with 3,072 cells. The side walls of the channel use symmetry boundary conditions. The inflow and outflow faces and the upper wall are set to far-field conditions. The bottom wall is deformed with a sinusoidal bump to create a reasonable variation in the flow, which has solid wall boundary condition.

Both the flow solution and the adjoint solutions are converged to a tolerance of $\mathcal{O}(10^{-12})$. We verify the adjoint accuracy with complex-step derivative approach given by: [6],

$$\frac{dF}{dx} = \frac{\text{Im}[F(x + ih)]}{h}, \quad (4.1)$$

where h is the complex step length. An imaginary step of $10^{-40}i$ is chosen as the perturbation.

Euler, Laminar NS, and RANS with both frozen-turbulence and full-turbulence are benchmarked against the complex step method. We choose Mach number 0.8




Figure 4.1: Volume mesh for bump verification case




Figure 4.2: C_p distribution of the RANS solution

Derivatives	Complex-Step	Adjoint	Difference
dC_D/dM	0.6529890 53	0.6529890 64	1.5E-8
dC_L/dM	1.6785453 80	1.6785453 72	4.9E-9
dC_D/dx	0.152323 769	0.152323 071	9.6E-8
dC_L/dx	0.0113249 74	0.0113249 75	4.6E-6

Table 4.1: Accuracy validations of the Euler adjoint

Derivatives	Complex-Step	Adjoint	Difference
dC_D/dM	0.6559854 01	0.6559854 67	1.0E-7
dC_L/dM	1.819804 777	1.819804 889	6.1E-8
dC_D/dx	0.011845 928	0.011845 836	7.7E-6
dC_L/dx	0.1453 07150	0.1453 12443	3.6E-5

Table 4.2: Accuracy validations of the Laminar NS adjoint

and Reynolds number 10 million for the flow condition. Figure 4.2 shows the C_p distribution of the RANS solution. Two objective functions, C_D and C_L , are used for verification. For the design variables, we choose Mach number to verify the aerodynamic derivatives and a point on the surface to verify the spatial derivatives. The results are summarized in Table 4.1 to Table 4.4.

We can see that the resulting derivatives match with complex-step solutions. The full-turbulence aerodynamic derivatives matches significantly better than the frozen-turbulence ones. Due to the complexity of the wall distance function in SA turbulence model, the wall distance computation is not linearized and is assumed constant in the turbulence model to simplify the automatic differentiation. Therefore, we see that the spatial derivatives have less accuracy than the aerodynamics derivatives for the full-turbulence adjoint.

4.2 Aerodynamic Shape Optimization of an ONERA M6 Wing

To demonstrate the effectiveness of the RANS adjoint formulation for aerodynamic shape optimization, an example of lift constrained drag minimization of a transonic

Derivatives	Complex-Step	Frozen-Turbulence Adjoint	Difference
dC_D/dM	0.673453 841	0.673684 112	3.4E-4
dC_L/dM	1.7 67928150	1.7 72398147	2.5E-3
dC_D/dx	0.009952 556	0.009952 686	1.3E-5
dC_L/dx	0.1 29946365	0.1 30232663	2.2E-3

Table 4.3: Accuracy validations of the frozen-turbulence adjoint

Derivatives	Complex-Step	Full-Turbulence Adjoint	Difference
dC_D/dM	0.673453841	0.673453842	1.1E-9
dC_L/dM	1.767928150	1.767928153	1.4E-9
dC_D/dx	0.009952556	0.009949493	3.1E-4
dC_L/dx	0.129946365	0.129890985	4.2E-4

Table 4.4: Accuracy validations of the full-turbulence adjoint

wing is presented. The particular test considered is the well-studied ONERA M6 wing [92]. This geometry has been studied by numerous authors [93, 94, 95, 96, 97, 10, 12] due to the simple, well defined geometry and the availability of experimental data.

The optimization problem considered is described below:

$$\begin{aligned}
& \underset{x}{\text{minimize}} && C_D(x) \\
& \text{subject to} && C_L \geq C_L^* \\
& && V \geq V_0 \\
& && t_i \geq 1, \quad i = 1, \dots, 21.
\end{aligned}$$

The objective is to reduce the drag coefficient while maintaining a specified lift coefficient, $C_L^* = 0.271$. The lift coefficient is based on a reference area of 0.75296 m^2 . Additional geometric constraints in the form of volume and thickness constraints are also used and are described in section 4.2.2.

4.2.1 Verification and Grid Refinement Study

Before optimizations were carried out, a grid refinement study and comparison with experimental data was made. The external flow condition for the experimental data and subsequent optimizations is:

$$M = 0.8395 \quad Re = 11.72 \times 10^6 \quad \alpha = 3.06^\circ \quad (4.2)$$

A sequence of 4 uniformly refined grids, labeled L1 through L4, were generated with grid sizes ranging from 129 thousand cells to over 66 million cells. The grids are generated using an in-house 3D hyperbolic mesh generator. The L2, L3 and L4 grids are all computed directly from their respective surface meshes while the L1 grid is obtained from the L2 grid by removing every other mesh node. An additional algebraic C-O topology Euler mesh was also generated for the purposes of comparing optimization results obtained with Euler and RANS analysis methods. The Euler

grid has approximately the same number of cells as the L2 RANS mesh to facilitate comparison between the computational cost for roughly equivalent Euler and RANS optimizations. A description of all grids used in this work is given in Table 4.5. For all grids the far-field surface is located approximately 100 Mean Aerodynamic Chords away from the body.


Grid	Cells	Surface Cells	Off-wall Cells	Off-wall Spacing	y_{\max}^+
RANS L1	129 024	4 032	32	3.0×10^{-6}	1.50
RANS L2	1 032 192	16 128	64	1.5×10^{-6}	0.67
RANS L3	8 257 536	54 512	128	0.75×10^{-6}	0.35
RANS L4	66 060 288	258 048	256	0.375×10^{-6}	0.18
Euler	1 044 480	18 432	40	3.0×10^{-4}	–

Table 4.5: ONERA M6 wing mesh sizes


The comparison of the experimental data with each of the four RANS grids is given in Figure 4.3. Overall, the flow solver has fairly accurately predicted the coefficient of pressure at each span-wise section. As expected, the finer grid resolutions do a better job of resolving both the location and strength of the shocks, although there is little discernible difference between the L3 and L4 grids. We believe the discrepancy between the computed and experimental data near the root can be attributed to wind tunnel effects and the splitter plate used in the physical setup that are not modeled computationally. A second discrepancy appears at the $2z/b = 0.90$ section where it is clear the position of the leading edge shock is displaced rearward as compared with the experimental data. This computational behavior is however, consistent with other results obtained on highly refined grids [98]. A possible explanation is due to small aeroelastic deformation of the physical model which not present in the computational model.

Additionally, an angle of attack sweep from 0° to 5° was run for each grid level to generate drag polars at the design Mach number of $M = 0.8395$. The polar is shown in Figure 4.4(a). It is clear that the coarsest grid, L1, is not sufficiently resolved for accurate drag prediction. Conversely, the L3 and L4 grids are nearly indistinguishable from each other except at the higher lift coefficients. While the discrepancy between the L2 and L4 grids is clearly visible it is fairly small and this level of refinement offers significantly computational savings compared to the L3 and L4 grids for the purposes of optimization.


Drag convergence curves for various angles of attack are given in Figure 4.4(b). The x -axis scale is given in terms of the *Grid Factor* which is defined as $N_{\text{cell}}^{-2/3}$. In general, the total drag coefficient decreases with increasing grid size. However,




(a) $2z/b = 0.20$




(b) $2z/b = 0.44$




(c) $2z/b = 0.65$



(d) $2z/b = 0.80$



(e) $2z/b = 0.90$



(f) $2z/b = 0.95$

Figure 4.3: C_p contours for each grid refinement level compared with experimental data

between the L3 and L4 grids at higher angles of attack, the trend reverses and the larger grids see a slight increase in drag. The root cause of this behavior is not known and warrants further investigation.

4.2.2 Geometric Parametrization, Constraints and Grid Movement


The geometric manipulation of the initial geometry is carried out using the Free Form Deformation (FFD) volume approach [62]. The design variables, x , are used to perturb the control points on a 3-dimensional parametric B-spline volume which in turn, perturbs the coordinates of the CFD surface mesh embedded parametrically inside. The design variable vector consists of 6 twist values that twist each of the six span-wise planes of control points, and 144 shape variables. Each shape variable perturbs individual coordinates of the FFD in the y direction. Note that since the root twist is allowed to vary, angle of attack is *not* a variable and the optimizations are carried out a fixed angle of attack of 3.06° .

To ensure a well-posed optimization problem, several additional geometric constraints are also employed. The internal volume of the wing is constrained to be greater than or equal to its initial value. A total of 21 thickness constraints are used; 10 distributed along the 15% chord line, 10 distributed along the 99% chord line and a single additional constraint near the mid-chord position at the wingtip. The leading edge constraints prevent a sharp leading edge from forming and the trailing edge constraints prevent a reduction in the thickness of the trailing edge.

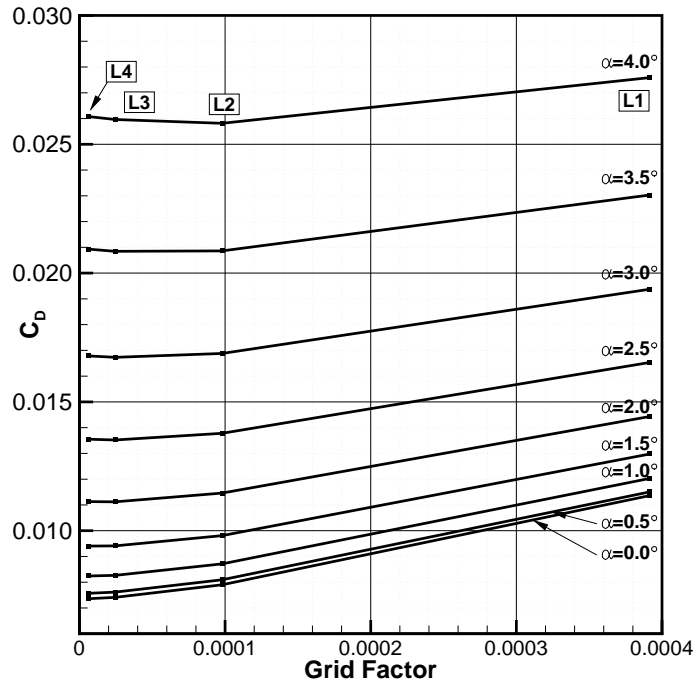
A view of the initial wing geometry, the FFD volume box and the thickness constraints are given in Figure 4.5. Note that the distribution of control points on the FFD are not uniform in the chord-wise direction. This clustering around the leading edge was used to ensure the optimizer is given sufficient geometric freedom to eliminate the leading edge shock present on the baseline design. Further, the blunt trailing edge of physical model is retained for the RANS simulations. A sharp trailing edge modification is used for the Euler grid.

The grids are deformed using a hybrid linear-elasticity algebraic mesh deformation algorithm previously developed by the Kenway [62] as described in Chapter 2. The mesh sensitivities required for the $\psi^T \frac{\partial A}{\partial x}$ computation are computed using Reverse Mode AD and a mesh adjoint equation.

A view of the surface mesh, symmetry plane and flow solution for the Euler and RANS grids are given in Figure 4.6.



(a) Drag polars for four levels of uniform refinement



(b) Grid convergence for C_D for various angles of attack.

Figure 4.4: Polar and grid convergence for each grid level.





Figure 4.5: FFD control points (blue spheres) and thickness constraints (red lines).

4.2.3 Optimization Algorithm


Due to high computational cost of the CFD solver, it is critical to choose an efficient optimization algorithm that requires a reasonably low number of function calls. Gradient-free methods, such as genetic algorithms, have a higher probability of getting close to the global minimum for cases with multiple local minima. However, slow convergence and large number of function calls would make gradient-free aerodynamic shape optimization infeasible with current computational resources. Therefore, we use gradient-based optimizers combined with adjoint gradient evaluations to achieve an efficient optimization process. For large number of design variables, the use of gradient-based optimizers is advantageous. We use a Python-based optimization package, pyOpt [72], to interface with CFD and adjoint solvers. We choose a gradient-based optimization algorithm, Sparse Nonlinear OPTimizer (SNOPT) [71], as the optimizer. SNOPT is a sequential quadratic programming (SQP) method, designed for large-scale nonlinear optimization problems with thousands of constraints and design variables. It uses a smooth augmented Lagrangian merit function and the Hessian of the Lagrangian is approximated using a limited-memory quasi-Newton method.

4.2.4 Computational Resources

The three optimizations are performed on a massively parallel supercomputer. Different processor counts are chosen for the Euler and RANS optimizations in an



(a) Euler grid



(b) RANS grid

Figure 4.6: Computational grids used for Euler and RANS analysis. C_p contours are shown for $M = 0.8395$ and $\alpha = 3.06^\circ$.

effort to keep the wall time of each optimization within a one day turn-around. Due to the lower computational and memory requirements for the Euler analysis, this optimization uses 32 processors while the two RANS optimizations use 88 processors.

4.2.5 Optimization Results

Three optimizations are considered: a RANS optimization employing the frozen turbulence assumption for the adjoint, a RANS optimization with the turbulence model linearization and an Euler optimization. An effort is made to compare the computational cost and accuracy of these differing approaches.

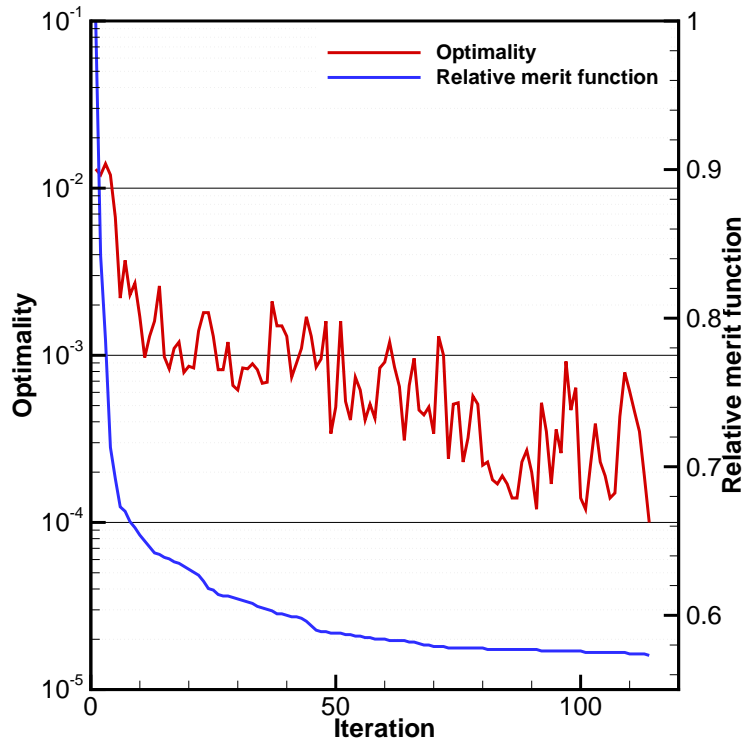
Firstly, we examine the convergence history of the optimizations, given in Figure 4.7. All optimization are converged to an optimality tolerance of 1×10^{-4} and take approximately 112 major iterations to reach this level of convergence.

Qualitatively, the merit function convergence for each optimization is similar: There is a very rapid decrease in C_D at the beginning of the optimization followed by much slower decreases as the optimization progresses. The first phase of the optimization involves the weakening of the two upper surface shocks. Referring to Figure 4.8(a), by the 10th iteration, the shocks have been entirely smoothed due to shape changes and this is responsible for the majority of the drag reduction. The second phase involves minor adjustments to the shape and modifications to the twist distribution. During this phase, an increase skin friction drag is traded for lower pressure and an overall decrease in the objective function. It is clear from Figure 4.8(b), that the majority of the wing twist present in the optimized design is added towards the end of the optimization, which is use primarily to reduce the induced drag of the wing.

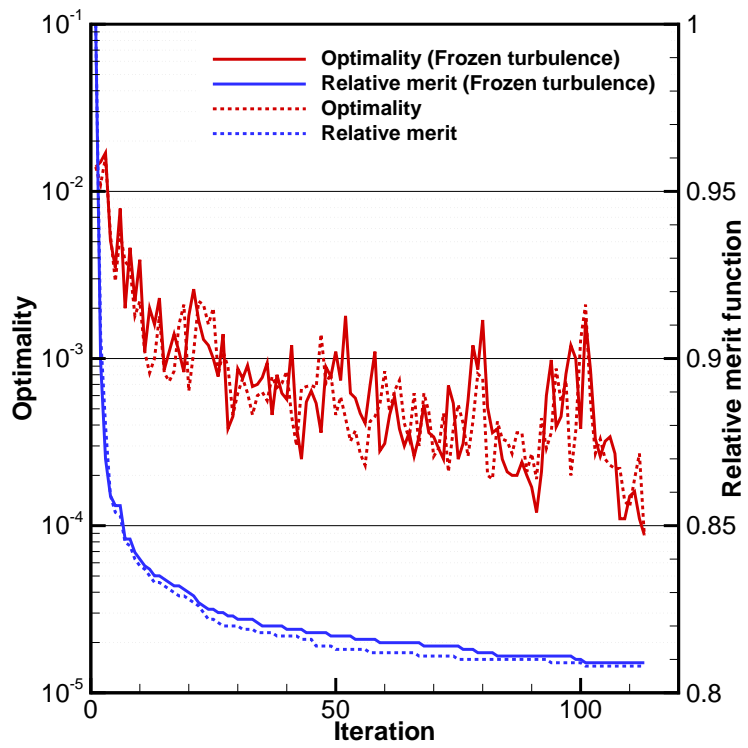
We now examine the cross sectional C_p contours of each of the three optimized designs. The same six span-wise locations as used in the experimental verification are reused. Figure 4.9 shows the contours for the baseline design, the frozen turbulence RANS optimization and the full RANS optimization. Figure 4.10 shows the baseline design, the optimized design and the optimized Euler design analyzed using RANS analysis. For this last case, the geometric design variables from the Euler optimization were used to perturb the L2 RANS grid and then obtain a solution at C_L^* .

Generally, the C_p contours for the two RANS optimization are similar. However, there are some slight differences, with the full RANS design resulting in somewhat smoother C_p contours. The largest discrepancy is observed on the lower surface near the leading edge.

A breakdown of the pressure and skin friction drag components is given in Ta-




(a) Convergence for Euler optimization




(b) Convergence for RANS optimization

Figure 4.7: Convergence history of optimality and relative merit function.

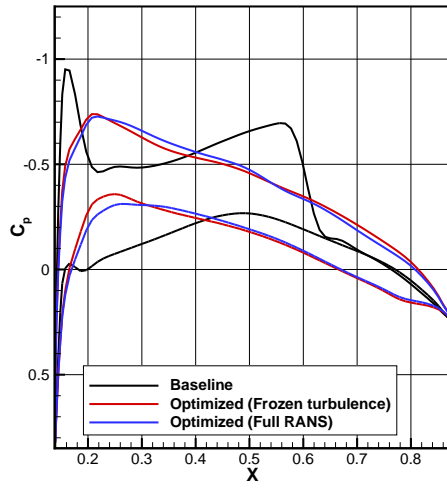


(a) Evolution of C_p during the optimization




(b) Evolution of cross section during the optimization


Figure 4.8: Evolution of C_p and cross section shape for section $2z/b = 0.65$.




(a) $2z/b = 0.20$




(b) $2z/b = 0.44$




(c) $2z/b = 0.65$



(d) $2z/b = 0.80$




(e) $2z/b = 0.90$




(f) $2z/b = 0.95$


Figure 4.9: C_p contours for RANS optimized designs.




(a) $2z/b = 0.20$




(b) $2z/b = 0.44$




(c) $2z/b = 0.65$



(d) $2z/b = 0.80$




(e) $2z/b = 0.90$




(f) $2z/b = 0.95$


Figure 4.10: C_p contours for Euler optimized design.




(a) $2z/b = 0.20$




(b) $2z/b = 0.44$




(c) $2z/b = 0.65$



(d) $2z/b = 0.80$



(e) $2z/b = 0.90$



(f) $2z/b = 0.95$

Figure 4.11: Cross section shapes for RANS and Euler optimized designs.

ble 4.6. In addition to the drag from the optimization (Optimized (L2)), we also analyze the baseline design and optimized design using the L3 grid. The goal is to verify that the gains made during the optimization are realized on the finer grid. This is indeed the case; The total drag reduction on the L3 grid is nearly identical to that on the L2 grid, justifying our choice of the L2 grid for optimization. For comparison, we also analyze the Euler design using the L2 RANS grid. The FFD approach greatly facilitates this exercise since the geometric design variables operate independently of the underlying mesh or surface topologies. Interestingly, the Euler optimized design shows remarkably good RANS performance, with the total drag coefficient only 3.7 counts higher than the RANS optimized design. Nevertheless, it worthwhile noting that this drag level was obtained by the RANS optimization after only 20 iterations, and corresponds to the initial optimization phase described previously. Most of the improvements small detailed shape and twist changes in the Euler optimization are evidently not realized in the viscous flow case.

Geometry	C_L	$C_{D_{total}}$	$\Delta C_{D_{total}}$	$C_{D_{pressure}}$	$\Delta C_{D_{pressure}}$	$C_{D_{friction}}$	$\Delta C_{D_{friction}}$
Baseline (L2)	0.2710	0.01725	–	0.01199	–	0.00526	–
Optimized (L2)	0.2710	0.01400	–0.00325	0.00847	–0.00343	0.00553	0.00027
Euler Design (L2)	0.2710	0.01437	–0.00288	0.00875	–0.00324	0.00561	0.00035
Baseline (L3)	0.2710	0.01687	–	0.01158	–	0.00529	–
Optimized (L3)	0.2710	0.01364	–0.00323	0.00816	–0.00342	0.00548	0.00019

Table 4.6: Drag break down for baseline and optimized designs on two mesh levels

A timing beak-down of the various components of each optimization is given in Table 4.7. The *Miscellaneous* category accounts for the time required for initial setup time, I/O, geometric manipulation, total sensitivity calculations and the optimization algorithm. Since the Euler optimization used few processors, the *Processor Hours* row indicates the true computational cost of the respective optimization. While, the full RANS simulation converges to a slightly better optimum in the same number iterations, the computation cost of the full RANS optimization is significantly more. Referring to Table 4.7, the main increase in cost for the full RANS simulation is the increased cost of residual assembly due to the extra state to be perturbed and the additional cost of solving the adjoint system. The increase in the adjoint solving cost is twofold: The matrix-vector products and preconditioner application is more costly due to the larger number of non-zeros (a factor of $(6/5)^2 = 1.44$) and the systems require more GMRES iterations for convergence. This results in the Full RANS

optimization requiring approximately 70% more CPU time than the frozen-turbulence assumption optimization. For the remainder of this section, for comparison, purposes we use the frozen turbulence optimization results.

Component	Euler		RANS (Frozen Turbulence)		RANS	
	Time (h)	Fraction	Time (h)	Fraction	Time (h)	Fraction
Flow solution	1.82	0.167	5.98	0.397	5.43	0.212
Adjoint assembly	0.70	0.065	1.32	0.088	1.73	0.067
Adjoint solution	7.88	0.724	7.37	0.490	17.93	0.701
Miscellaneous	0.49	0.045	0.37	0.025	0.50	0.020
Wall-time Total	10.89	1.000	15.04	1.000	25.59	1.000
Processor Hours	348.5	–	1323.52	–	2251.0	–

Table 4.7: Timing breakdown for each optimization

Origin	Grid Size	N _{DV}	Initial			Optimized			
			C _L	C _D	L/D	C _L	C _D	ΔC_D (%)	L/D
Present Work (L3)	8M	150	0.2710	0.01687	16.06	0.2710	0.01364	–19.1	19.86
Present Work (L2)	1M	150	0.2710	0.01725	15.71	0.2710	0.01400	–18.8	19.36
Osusky and Zingg [12]	2.2M	226	0.2590	–	–	0.2590	–	–17.1	–
Bueno-Orovio et al. [99]	43K	12	–	0.01712	–	–	0.01558	–10.0	–
Le Moigne and Qin [96]	312K	86	0.2697	0.01736	15.54	0.2964	0.01478	–14.9	18.23
Neilson and Anderson [95]	359K	21	0.2530	0.01680	15.06	0.2530	0.01420	–15.5	17.82
Lee et al. [97]	291K	40	0.2622	0.01751	14.97	0.2580	0.01586	–9.4	16.27

Table 4.8: Comparison of aerodynamic coefficients with previous work.

We compared our optimization results with previous optimization studies of the ONERA M6 wing, summarized in Table 4.8. We obtained a large drag reduction of 19.1%. Present work used a relatively large grid size and comparatively large number of design variables. Due to the clustering of shape design variables near the leading edge, the optimization was able to completely eliminate the leading edge shock.


The C_p contours for both the Euler and RANS optimized designs are shown in Figure 4.12. We can see that both Euler and RANS achieved a shock-free solution. The Euler optimized design has a rapid pressure recovery near the TE. The RANS optimized solution, however, has parallel pressure contour lines with nearly constant spacing, indicating a gradual increase of pressure.

We also investigated the drag divergence at different C_L . Figure 4.13 shows the drag divergence plot for three different C_L values. Both Euler and RANS optimized design reduced drag over the entire Mach range and the divergence Mach number are increased at all C_L values as compared to the baseline design. The drag coefficient remains nearly constant up to the divergence Mach numbers. At higher C_L , we see a drag pocket at the optimized Mach number for Euler solution. However, the drag dip on the RANS design is not significant. The effect could be due to the relatively low C_L . The drag dip at the optimized Mach number may become more prominent at higher loadings.


The baseline designs have lift distributions that are already reasonably close to elliptic. Both Euler and RANS optimized designs result in lift distributions that are very close to the optimum elliptical distribution, as shown in Figure 4.14. As a result, lift-induced drags were decreased, contributing to the pressure drag reduction shown in Table 4.6. The shift in lift distributions were obtained by the change in twist distributions shown in Figure 4.15. We also see that the Euler optimization tends to change t/c more significantly than the RANS optimization.

4.3 Conclusions

In this chapter, we verified the aerodynamic and geometric derivatives computed with adjoint method against complex-step method using a flow over a bump test case. The derivative computations are accurate, robust and efficient. The aerodynamic gradients differ by $\mathcal{O}(10^{-9})$, and the spatial gradients differ by $\mathcal{O}(10^{-4})$ when compared with the complex-step method. A RANS aerodynamic shape optimization of the ONERA M6 wing is presented as a preliminary test case. The results are compared with a design obtained by a comparable Euler optimization. We achieved




(a) Euler optimization




(b) RANS optimization

Figure 4.12: C_p contours for baseline and optimized designs for Euler and RANS

a drag reduction of 19% as compared the baseline wing. The shocks on the upper surface was completely eliminated and the optimized design improved the drag coefficient at all flight Mach numbers. The drag divergence Mach number of the optimized design is also increased. For the ONERA M6 optimization problem considered, the full RANS adjoint formulation resulted in a slightly better optimized design, but the optimization was 70% more costly than the frozen turbulence formulation. With this verification of the aerodynamic shape optimization framework, we are now ready to tackle practical aerodynamic design optimization problems in the following chapters.



(a) Euler optimization



(b) RANS optimization

Figure 4.13: Drag divergence curves for three fixed lift coefficients.




Figure 4.14: Lift distribution for baseline and optimized designs.




Figure 4.15: Thickness-to-chord ratio and twist distributions for baseline and optimized designs.

CHAPTER 5

Gradient-Based and Gradient-Free Optimization

Aerodynamic design optimization requires large computational resources, since each design evaluation requires the solution of a system of partial differential equations in a three dimensional domain. Thus, the choice of optimization algorithm is critical, as it directly affects the number of required design evaluations to reach the optimum design. To help designers make an informed choice, we benchmark several optimization algorithms in this thesis, including gradient-based and gradient-free methods using three test problems of increasing difficulty: a multi-dimensional Rosenbrock function, a RANS-based aerodynamic twist optimization problem and an aerodynamic shape optimization problem. The majority of the gradient-based optimizers successfully solved all three test problems, while the gradient-free methods require two to three orders of magnitude more computational effort when compared to the gradient-based methods. Thus, gradient-based algorithms are the only viable option for solving large-scale aerodynamic design optimization problems.

The optimization can be performed with gradient-based or gradient-free methods. Gradient-based methods are best when an efficient gradient evaluation is available. The computational expense of evaluating the gradient using finite difference or complex step methods [6] is prohibitive for aerodynamic shape optimization with respect to hundreds of variables. The adjoint method, however, can provide accurate and efficient gradient evaluations [2, 28], and adjoint-based aerodynamic shape optimization has been widely used [29, 11, 100, 101, 102]. Gradient-free methods are generally simpler to implement, and claim to find the global optimum, but the computational cost is higher. In this chapter, we investigate the local optima in aerodynamic shape optimization of a transonic wing. In addition, we also compare the optimization algorithms using this benchmark.

The aerodynamic shape optimization has been well-studied using various approaches. Sasaki *et al.* [103] applied an adaptive range multiobjective genetic al-

gorithm (ARMOGA) to aerodynamic wing design. A four-objective optimization of wing shape and planform were presented using 72 design variables, subject to thickness and planform shape constraints. Moigne and Qin [96] studied aerodynamic shape optimization based on a discrete adjoint of a Reynolds-averaged Navier–Stokes (RANS) solver. A variable-fidelity optimization method combining low- and high-fidelity models was used. The optimization reduced 23% drag on a RAE2822 airfoil and 15% on a ONERA M6 wing. Their results showed that using a variable-fidelity method that performs most of the optimization on a low-fidelity, low-cost model (Euler equations on a coarse grid) reduces the overall computing time.

Lyu *et al.* [101] presented the results of lift-constrained drag minimization of the AIAA Aerodynamic Design Optimization Discussion Group (ADODG) Common Research Model wing using a RANS solver. A 8.5% drag reduction was achieved using a multilevel optimization approach. The same optimization was also performed starting from a randomly generated initial design, and closely spaced local optima were observed.

Several authors compared the performance of different optimization methods. Foster and Dulikravich [104] compared a hybrid gradient method and a hybrid genetic algorithm for a three dimensional aerodynamic lifting body design. Zingg *et al.* [105] performed a comparison of genetic algorithm and gradient methods in aerodynamics airfoil optimization. Genetic algorithm required 5–200 times more function evaluations than gradient-based methods with adjoint sensitivity. They suggested genetic algorithm was more suited for preliminary design with low-fidelity models. Gradient-based optimizers may be more appropriate for detailed designs with high-fidelity models. Obayashi and Tsukahara [106] compared a gradient-based method with simulated annealing, and a genetic algorithm on an airfoil lift maximization problem. The genetic algorithm required the highest number of function evaluation. However, the genetic algorithm achieved the best design compared to the other two methods. Frank and Shubin [107] compared one-dimensional duct flow optimization with finite-difference gradients, optimization with analytic gradients, and an all-at once method where the flow and design variables are simultaneously altered. They concluded that the optimization with analytic gradients was the best approach that can be retrofitted to most existing codes.

In this chapter, we extend the previous studies of optimizer comparison and local optima using high-fidelity aerodynamic shape optimization. We compare several optimization algorithms including 6 gradient-based methods—SNOPT, PSQP, SLSQP, IPOPT, CONMIN, GCMMA—and 2 gradient-free methods—ALPSO, NSGA2. We

test those optimizers using a multi-dimensional Rosenbrock function, a wing twist optimization problem, and a wing shape optimization problem. The strengths and weaknesses of each method are discussed. This work is based on a previous paper presented by the author. [108]

5.1 Multi-dimensional Rosenbrock Function

To examine the effectiveness of the optimizers listed above, we first minimize a multi-dimensional Rosenbrock function [109]. In addition, a nonlinear constraint is added to the formulation, and the complete problem is:

$$\begin{aligned} & \text{minimize} && \sum_{i=1}^{n-1} 100(x_{i+1} - x_i^2)^2 + (x_i - 1)^2 \\ & \text{with respect to} && x \in \mathbb{R}^n \\ & \text{subject to} && \sum_{i=1}^{n-1} (1.1 - (x_i - 2)^3 - x_{i+1}) \geq 0 \end{aligned}$$

The constraint is always active at the optimum. For a two-dimensional problem, the feasible optimum is at [1.2402, 1.5385] with an objective value of 0.0577244. The optimizations are started from $x_i = 4$, and the design variables are bounded so that they remain in the interval $[-5.12, 5.12]$.

We set the options for each optimizer based on our best knowledge. For example, we use a swarm size of 8 and a maximum outer iteration of 4000 for ALPSO. We use a population size of 24 and 200 generations for NSGA2. We terminate all optimizations with 10^{-6} relative tolerance of 3 consecutive iterations and 10^{-6} feasibility tolerance. In this study, we investigate the computational cost and effect of increasing number of design variables. In addition, we compare results found using finite-difference gradients with those found using analytical derivatives.

Figure 5.1 shows the optimization path taken by each optimizer. Gradient-based methods follow through the Rosenbrock valley and converge toward the optimum. Gradient-free methods converge their population toward the optimum in a more scattered way. The convergence history of selected optimizers of the two dimensional Rosenbrock function is plotted in Figure 5.2. For gradient-free methods, only the best point is plotted for each iteration or generation. Most of the gradient-based optimizers converge to a tolerance of 10^{-5} within 150 iterations, while ALPSO converges to the same tolerance using 3,368 iterations and NSGA2 can not converge to the same tolerance before we terminate the computation. NSGA2 terminates when

the maximum number of generation (200) is reached. SLSQP is the fastest, with 34 function evaluations.

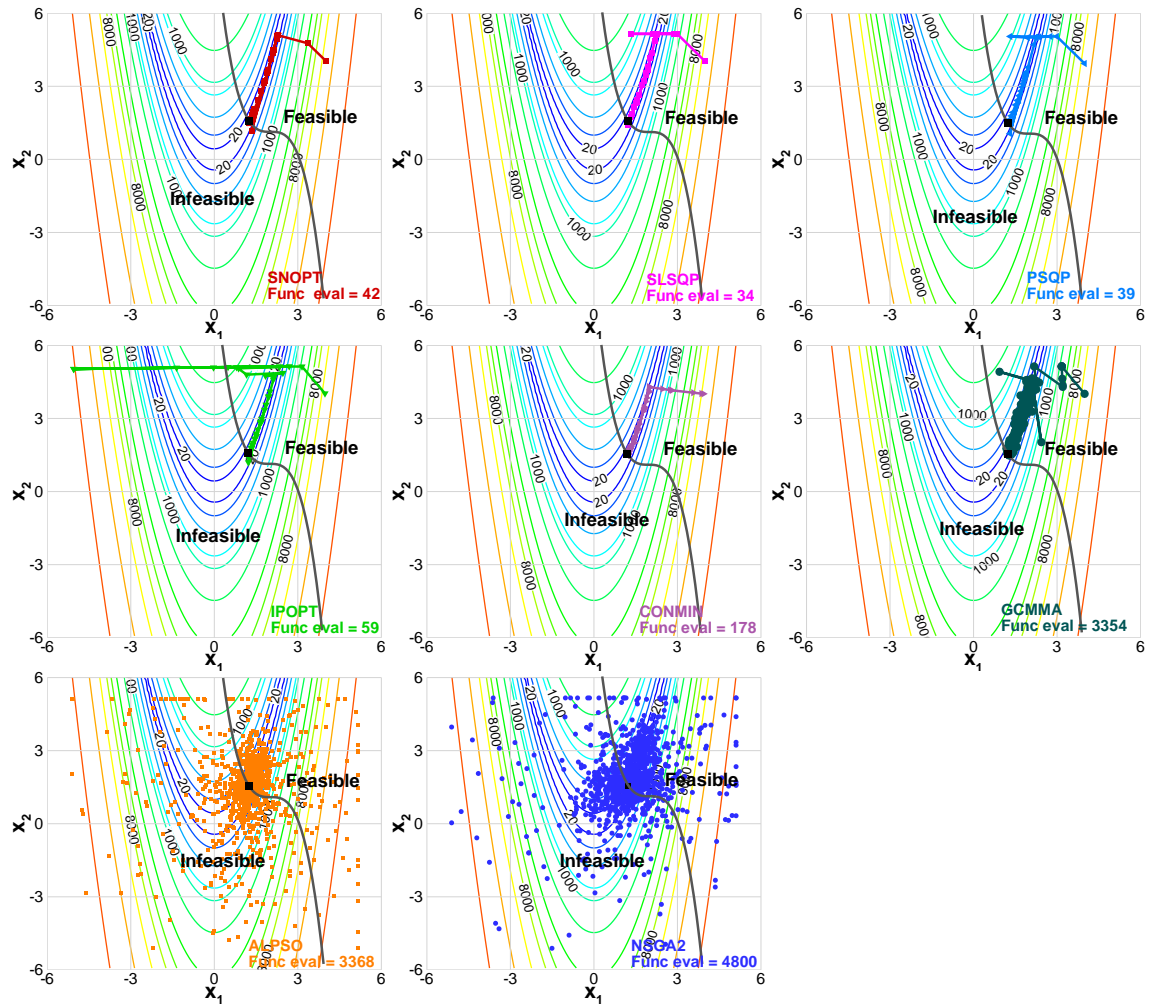


Figure 5.1: Optimization paths for the constrained 2D Rosenbrock function

To visualize the effect of increasing the dimensionality of the problem, we also plot the number of function evaluations required to converge the optimization for an increasing number of design variables. As shown in Figure 5.3, the gradient-free methods tend to have quadratic or cubic growth of function evaluations with increasing dimensionality, while the gradient-based methods follow a linear trend. The difference between gradient-based methods with finite-difference gradients and gradient-based methods with analytical gradients is significant, motivating the use of the adjoint method in aerodynamic shape optimization that we discuss later.

To investigate the local minima, we remove the constraint. Then, two minima (one local and one global) occur for higher dimensions [110]. The local minimum is at




Figure 5.2: Gradient-based methods converge faster for the Rosenbrock problem

$[-1, 1, 1, \dots]$, and the global minimum is at $[1, 1, 1, \dots]$. We start the optimization at $[-0.8, 0.8, 0.8, \dots]$, which is relatively close to the local minimum at $[-1, 1, 1, \dots]$. All optimizers were able to converge to the global minimum for 2 and 4 design variables. However, for 8 design variables or more, gradient-based methods converge to the local minimum, while the gradient-free methods find the global minimum, as shown in Figure 5.4.

In this study, we compare the optimizers using a multi-dimensional Rosenbrock function. Gradient-free methods take 2 to 4 orders of magnitude more function evaluations to converge the optimization than most gradient-based methods. NSGA2 cannot achieve the required accuracy within 200 generations. The gradient-free methods have a higher probability of converging to a point near a global optimum. However, it requires high number of function evaluations with large number of design variables, making it infeasible for large-scale aerodynamic shape optimization. Thus, gradient-based methods with efficient gradient computations are a better choice for large-scale




Figure 5.3: The gradient-free methods require an excessive number of function evaluations for large numbers of variables.

optimizations.

5.2 Aerodynamic Twist Optimization

In this study, the objective is to perform a lift-constrained drag minimization of the Common Research Model (CRM) wing [36, 37, 38]. The flow is solved using RANS equations. The adjoint method is used to solve the gradients including the linearization of the turbulence model. The baseline geometry is the same as the one used by Lyu *et al.* [101]. The specifications are given by the Aerodynamic Design Optimization Discussion Group. The mesh is generated using an O-grid topology, extruded to a farfield at a distance equal to 25 times the wing span. Grid size and y^+ are listed in Table 5.1. We use level 3 and level 2 grids in this study.

For this problem, we use 8 wing twist design variables to provide a reasonable run time to compare gradient-based and gradient-free optimizers. A lift coefficient




Figure 5.4: Visualization of the local and global minimum of 8-dimensional Rosenbrock function

Grid level	Grid size	y+
L2	450,560	2.213
L3	56,320	8.4086

Table 5.1: Grid size used in aerodynamic twist optimization

constraint of $C_L = 0.5$ is imposed. The initial wing has zero twist. The coarse L3 grid is used. We also perform the optimization on the L2 grid using gradient-based methods.

The optimized twist, lift and pressure distributions using each optimizer and the L3 grid are shown in Figure 5.5. All optimizers except NSGA2 converge to the same drag value. The difference is within 0.1 of a drag count, and the twist distributions

are nearly identical.

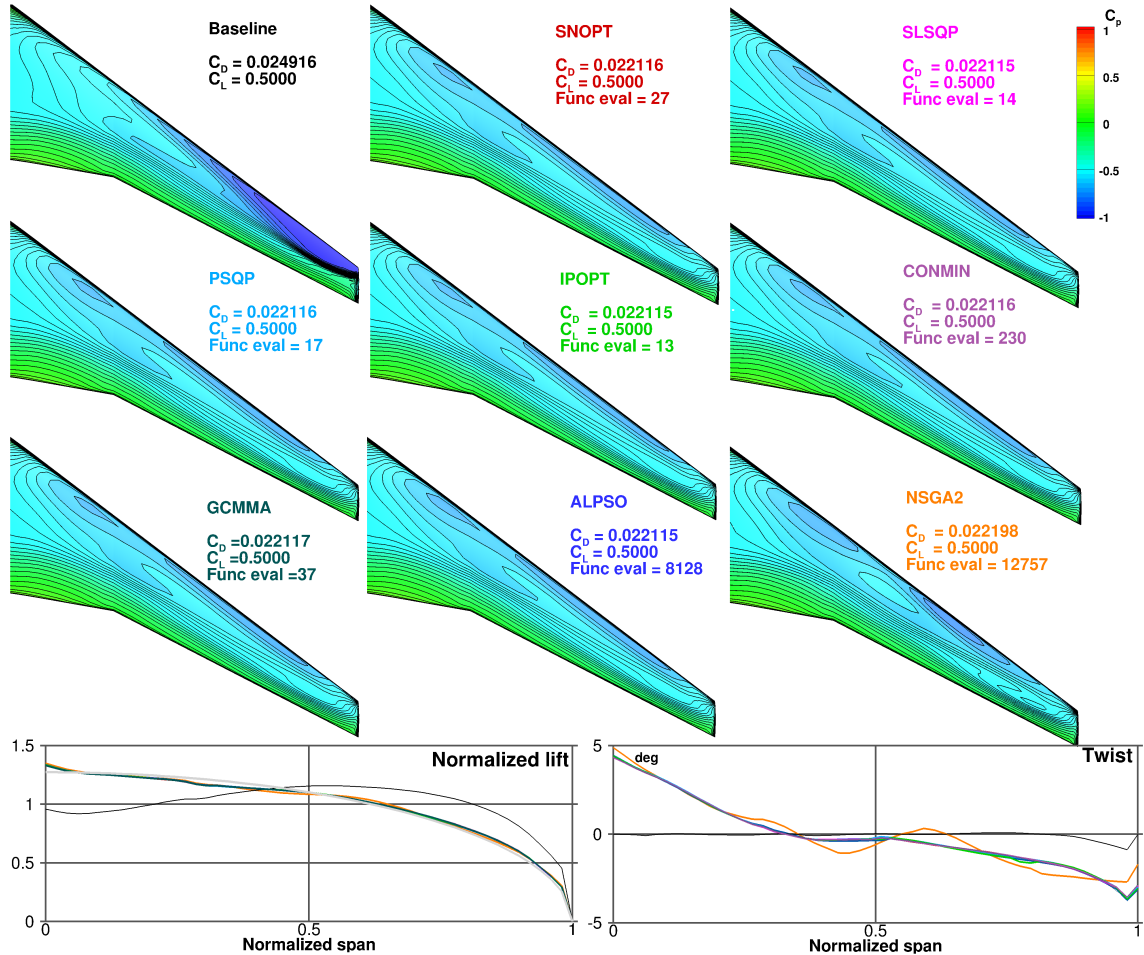


Figure 5.5: Aerodynamic twist optimization comparison on the L3 grid

The gradient-free optimizers take 3 orders of magnitude more iterations than the gradient-based optimizers. We compare the computational cost of the optimizers in Table 5.2. The relative convergence tolerances for gradient-based methods are 10^{-5} for the objective, and 10^{-4} for the constraints. The corresponding values for the ALPSO optimizer are 10^{-2} for the objective and 10^{-3} for the constraints. For this case, SLSQP, PSQP and IPOPT perform well. CONMIN is slower and did not achieve the required tolerance. For non-gradient methods, ALPSO performs better than NSGA2, as it takes half of the time of NSGA2, and converges to a better design. ALPSO converges to the same optimum as the gradient-based methods, while the optimum obtained by NSGA2 is 0.8 drag count higher with a different twist distribution.

The convergence history of the optimization is shown in Figure 5.6 and 5.7. Since the number of function evaluation for the gradient-based methods are two orders of

Optimizer	Iteration numbers	Proc hours
SNOPT	27	5.81
SLSQP	14	2.56
PSQP	17	3.70
IPOPT	13	2.69
CONMIN	230	33.61
GCMMA	37	4.57
ALPSO	8129	1695.72
NSGA2	12757	2744.16

Table 5.2: Computational cost comparison of the twist optimization for the L3 grid

magnitude lower than gradient-free methods, we plot the convergence of the gradient-based methods and the gradient-free methods separately. For gradient-based methods, we plot the value of objective function with respect to the number of function evaluations. For the gradient-free methods, only the best point of each iteration or generation is plotted.

After performing the comparison for the L3 grid, the same optimization is verified using the finer L2 grid. Figure 5.8 shows the optimized results of the gradient-based optimizers using the L2 grid. The L2 optimization is too costly to be implemented with gradient-free methods. Using only twist design variables, the shock on the wing can not be completely removed. The drag is reduced by 29 counts. Similarly, the difference in drag between the optimizers is within 0.1 count. Thus, it seems as if the twist optimization problem has only one optimum. The lift distributions of optimized design using L2 and L3 grids are shown in Figure 5.9. The difference in grid resolutions results in a difference in the optimized twist distribution. The optimized design increases lift at the root and reduces lift at the tip, thus moving towards an elliptical lift distribution. However, since the optimizers minimize the total drag with only 8 twist design variables, the optimal trade-off between induced drag with wave and viscous drag is not obvious, resulting in a non-elliptical lift distribution.

In this study, we examined a twist design problem. We used 8 design variables subject to a lift constraint. We compared the optimized results using different optimizers on two grid levels. All optimizers converged to a similar optimum. A single global minimum is observed. The gradient-based methods converged significantly faster than the gradient-free methods.




Figure 5.6: Twist optimization convergence history for gradient-free methods for the L3 grid




Figure 5.7: Twist optimization convergence history for gradient-based methods for the L3 grid

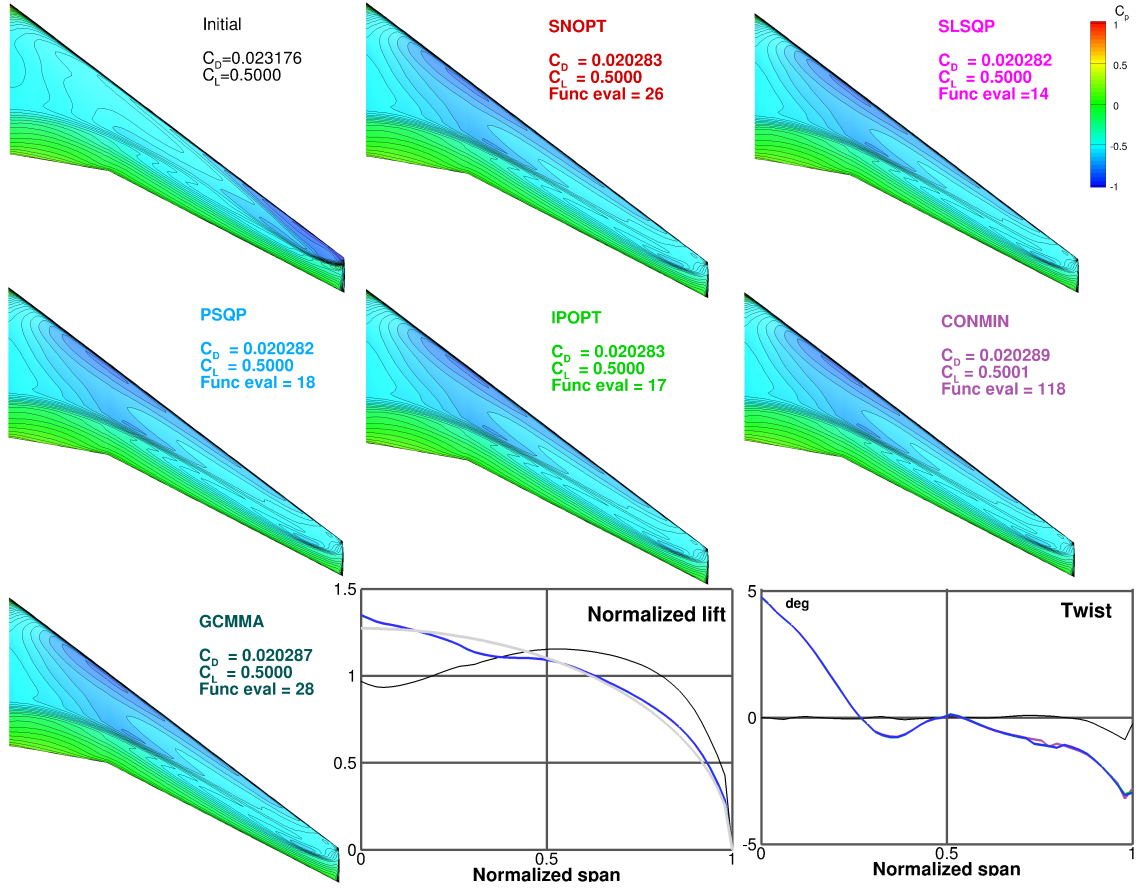


Figure 5.8: Aerodynamic twist optimization comparison for the L2 grid

5.3 Aerodynamic Shape Optimization

In this study, we use the same geometry as the twist optimization case discussed above. Instead of using just 8 twist design variables, a total of 192 shape design variables are considered. As in the previous case, the angle-of-attack is also allowed to vary, and we perform drag minimization subject to a lift constraint of $C_L = 0.5$. The wing thickness is constrained from reducing relative to the initial geometry by imposing 750 thickness constraints. In addition, a volume constraint is imposed to ensure that the internal volume does not decrease beyond the baseline volume. This problem requires significantly more computational resources than the previous case. We perform the shape optimization using 4 different gradient-based optimizers on the L2 grid. The convergence tolerance is 10^{-6} for the objective and 10^{-4} for the constraints.

Figure 5.10 shows the final design resulting from the use of different optimizers. The results from the baseline wing are shown in black. More detailed comparisons for




Figure 5.9: Lift distribution comparison of the twist optimized designs

each optimizer are shown in Figures A.1– A.4. The drag is reduced by 4.84%, from 206.7 to 196.6 counts, which is similar to the previous result [101].

We can see that all optimizers achieve a shock-free wing with an elliptical lift distribution. The baseline design has a strong shock, as evidenced by closely spaced C_p contours, while the optimized designs have a parallel, equally spaced pressure contours. The variation in C_D is within 1 drag count between the various optimizers. All optimized shapes are similar to each other, and only small difference in shape are observed. The comparison of the computational time for various optimizers is shown in Table 5.3. SNOPT converges the fastest among all optimizers. The optimized results using GCMMA is 0.2 drag count higher than the others. The convergence history is plotted in Figure 5.11.

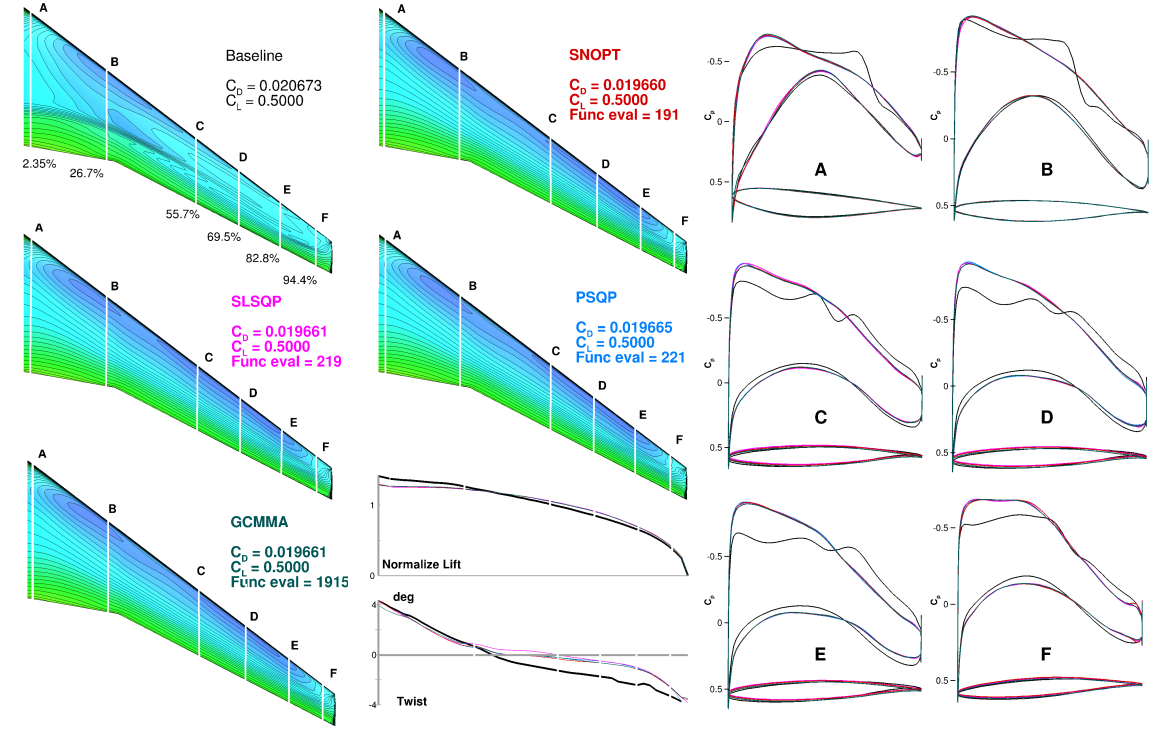


Figure 5.10: Aerodynamic shape optimization comparison for the L2 grid

Optimizer	Function evaluations	Proc hour
SNOPT	92	224.98
SLSQP	116	306.38
PSQP	221	562.60
GCMMA	298	772.60

Table 5.3: Computational cost comparison of the shape optimization for the L2 grid

5.4 Conclusions

We evaluated several optimization algorithms for three different aerodynamic shape optimization problems. The algorithms we considered included gradient-based methods with adjoint gradients and gradient-free methods (a particle swarm optimization and a genetic algorithm). The gradient-free methods required 2 to 4 orders of magnitude more iterations than gradient-based methods. We conclude that gradient-based methods with adjoint gradients are the best choice for solving large-scale aerodynamic design optimization problems.




Figure 5.11: Shape optimization convergence history for gradient-based methods for the L2 grid

CHAPTER 6

Aerodynamic Design Optimization Acceleration Methods

Aerodynamic shape optimization based on high-fidelity models is a computational intensive endeavor. The majority of the computational time is spent in the flow solver, and in the gradient calculation. In this chapter, we present two approaches for reducing the overall computational cost of the optimization. The techniques are tested using the Common Research Model wing benchmark defined by the Aerodynamic Design Optimization Discussion Group (ADODG). The aerodynamic model solves the Reynolds-averaged Navier–Stokes equations with a Spalart–Allmaras turbulence model. A gradient-based optimization algorithm is used in conjunction with an adjoint method that computes the required derivatives. The drag coefficient is minimized subject to lift, pitching moment, and geometric constraints.

The majority of the computational time is spent in the flow solver and in the gradient calculation. There are several possible ways to reduce the overall optimization time. One way is to reduce the flow solution time. This has been extensively researched by the CFD community. Commonly used methods, such as multigrid [111], pre-conditioning [112], and variations on Newton-type methods [113, 114], can improve the convergence of the solver, thus reducing the overall optimization time.

The second way to reduce the overall computational cost is to reduce the gradient computational time, which was pioneered by Jameson [2] through the development of adjoint method, which efficiently compute gradients with respect to large numbers of shape design variables. With an efficient adjoint implementation, the cost of computing the gradient of a single function of interest with respect to hundreds or thousands of shape design variables is roughly of the same order of the cost of one flow solution [28]. Those methods have been successfully applied in recent aerodynamic shape optimizations [30, 32, 33, 100, 101, 102].

As these numerical techniques become mature, there is a need to research new ways to reduce the optimization cost. In this chapter, we explore two approaches to further improve the efficiency of aerodynamic shape optimization. The first idea is to use Richardson's extrapolation to approximate the flow solution and gradients of a fine grid using the results of coarse grids. The second approach is to perform grid sequencing at the optimization level. This is inspired by the multigrid method in CFD. We use smaller grids to accelerate the convergence of a large grid. We perform the optimization first on a smaller grid until a certain level of optimality is achieved. Then, we move on to the next grid level and use the design variables from the previous grid level as the initial design variables. This process is repeated until the last grid level has converged. We compare both approaches with a direct optimization with only the fine grid. The comparison of the optimized designs, convergence and computational cost is presented in the thesis.

This study is performed using the benchmark case for aerodynamic design optimization developed by the Aerodynamic Design Optimization Discussion Group (ADODG): the lift-constrained drag minimization of the NASA Common Research Model (CRM) wing [36, 37, 38] with a RANS model that were presented at the 2014 AIAA Science and Technology Forum and Exposition in a special session organized by the ADODG [115, 101, 116, 117, 118]. The work in this Chapter is based on a previous paper presented by the author. [119]

6.1 Problem Formulation

The optimization case we used in this study is to perform lift-constrained drag minimization of the NASA CRM wing using the RANS equations. In this section, we provide a complete description of the problem.

6.1.1 Baseline Geometry

The baseline geometry is a wing with a blunt trailing edge extracted from the CRM wing-body geometry [37, 38]. The NASA CRM geometry was developed for applied CFD validation studies. The CRM is representative of a contemporary transonic commercial transport, with a size similar to that of a Boeing 777. The CRM geometry has been optimized, but several design features, such as an aggressive pressure recovery in the outboard wing, were introduced into the design to make it more interesting for research purposes and to protect intellectual property. This baseline geometry provides a reasonable starting point for the optimization, while leaving room




Figure 6.1: Baseline CRM wing geometry scaled by its mean aerodynamic chord.

for further performance improvements. In addition, the CRM was designed together with the fuselage of the full CRM configuration, so its performance is degraded when only the wing is considered.

The geometry and specifications are given by the ADODG. The fuselage and tail are removed from the original CRM, and the root of the remaining wing is moved to the symmetry plane. This baseline geometry is shown in Fig. 6.1. All coordinates are scaled by the mean aerodynamic chord (275.8 in). The resulting reference chord is 1.0, and the half span is 3.758151. The moment reference point is at $(x, y, z) = (1.2077, 0.0, 0.007669)$, while the reference area is 3.407014.

6.1.2 Geometric Parametrization

We use a free-form deformation (FFD) volume approach to parametrize the wing geometry as described in Chapter 2. Figure 6.2 shows the FFD volume and the geometric control points used in the aerodynamic shape optimization. The shape design variables are the displacement of all FFD control points in the vertical (z) direction.




Figure 6.2: The shape design variables are the z -displacements of 720 FFD control points (red spheres).

6.1.3 Mesh Convergence Study

We generate the mesh for the CRM wing using an in-house hyperbolic mesh generator. The mesh is marched out from the surface mesh using an O-grid topology to a farfield located at a distance of 25 times the span (about 185 mean chords). The nominal cruise flow condition is Mach 0.85 with a Reynolds number of 5 million based on the mean aerodynamic chord. The mesh we generated for the test case optimization contains 3.6 million cells. Mesh size and aerodynamic coefficients under the nominal operating condition are listed in Table 6.1.

Mesh level	Mesh size	C_D	C_L	C_M	α
$h = 0$	∞	0.01990			
L00	230,686,720	0.01992	0.5000	-0.1776	2.2199 ^o
L0	28,835,840	0.01997	0.5000	-0.1790	2.2100 ^o
L1	3,604,480	0.02017	0.5000	-0.1810	2.1837 ^o
L2	450,560	0.02111	0.5000	-0.1822	2.1944 ^o

Table 6.1: Mesh convergence study for the baseline CRM wing.






Figure 6.3: The mesh convergence study shows that the difference between the drag values computed with the 28.8 M and the 230.7 M grids is within 1 count.


We perform a mesh convergence study to determine the resolution accuracy of this mesh. Table 6.1 lists the drag and moment coefficients for the baseline meshes. We also compute the zero-grid spacing drag using Richardson’s extrapolation, which estimates the drag value as the grid spacing approaches zero [120]. The zero-grid spacing drag coefficient is 199.0 counts for the baseline CRM wing. We can see that the L0 mesh has sufficient accuracy: the difference in the drag coefficient for the L0 mesh and the zero-grid spacing drag is within one drag count. The surface and symmetry plane meshes for the L0, L1, and L2 grid levels are shown in Fig. 6.4.



(a) L0 mesh: 28.8 M cells, 199.7 drag counts.



(b) L1 mesh: 3.6 M cells, 201.7 drag counts.



(c) L2 mesh: 450 k cells, 211.1 drag counts.

Figure 6.4: O-grids of varying sizes were generated using a hyperbolic mesh generator.

6.1.4 Optimization Problem Formulation

The aerodynamic shape optimization seeks to minimize the drag coefficient by varying the shape design variables subject to a lift constraint ($C_L = 0.5$), and a pitching moment constraint ($C_{My} \geq -0.17$). The shape design variables are the z -coordinate movements of 720 control points on the FFD volume (shown in Fig. 6.2) and the angle-of-attack. The control points at the trailing edge are constrained to avoid any movement of the trailing edge. Therefore, the twist about the trailing edge can be implicitly altered by the optimizer using the remaining degrees of freedom. The leading-edge control points at the wing root are also constrained to maintain a constant incidence for the root section. There are 750 thickness constraints imposed in a 25 chordwise and 30 spanwise grid covering the full span and from 1% to 99% local chord. The thickness is set to be greater than 25% of the baseline thickness at each location. Finally, the internal volume is constrained to be greater than or equal to the baseline volume. The complete optimization problem is described in Table 6.2.

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

Table 6.2: Aerodynamic shape optimization problem




Figure 6.5: Sensitivity study of the baseline wing shows which shape changes yield the largest improvements.

6.1.5 Surface Sensitivity on the Baseline Geometry

To examine the potential improvements of the baseline geometry, we performed a sensitivity analysis in [101]. The sensitivity of the drag and pitching moment with respect to the airfoil shape is shown in Fig. 6.5 as a contour plot of the derivatives of C_D and C_{M_y} with respect to shape variations in the z direction. The regions with the highest gradient of C_D are near the shock on the upper surface and near the wing leading edge. This indicates that leading-edge shaping and shock reduction through local shape changes should be the major drivers in C_D reduction at the beginning of the optimization. As for C_{M_y} , the shape changes near the root and tip of the wing are the most effective in adjusting the pitching moment. Since these sensitivity plots are a linearization about the current design point, they provide no information about

the constraints. Nonetheless, these sensitivity plots indicate what drives the design at this design point.

6.2 Aerodynamic Shape Optimization Results

Three different aerodynamic shape optimization approaches are presented: direct optimization, optimization with Richardson’s extrapolation, and optimization with a multilevel approach. The optimized designs and the efficiency are compared in this section.

6.2.1 Direct Aerodynamic Shape Optimization

We present the direct aerodynamic design optimization for the CRM wing benchmark problem (described in Table 6.2) under the nominal flight condition (Mach 0.85, $Re = 5 \times 10^6$). We use the L1 grid (3.6 M cells) directly for the optimization. This is currently the most common way to perform aerodynamic shape optimization due to its simplicity. The optimization is computed with 64 processors. Figure 6.6 shows a detailed comparison of the baseline wing and the optimized wing using direct optimization.

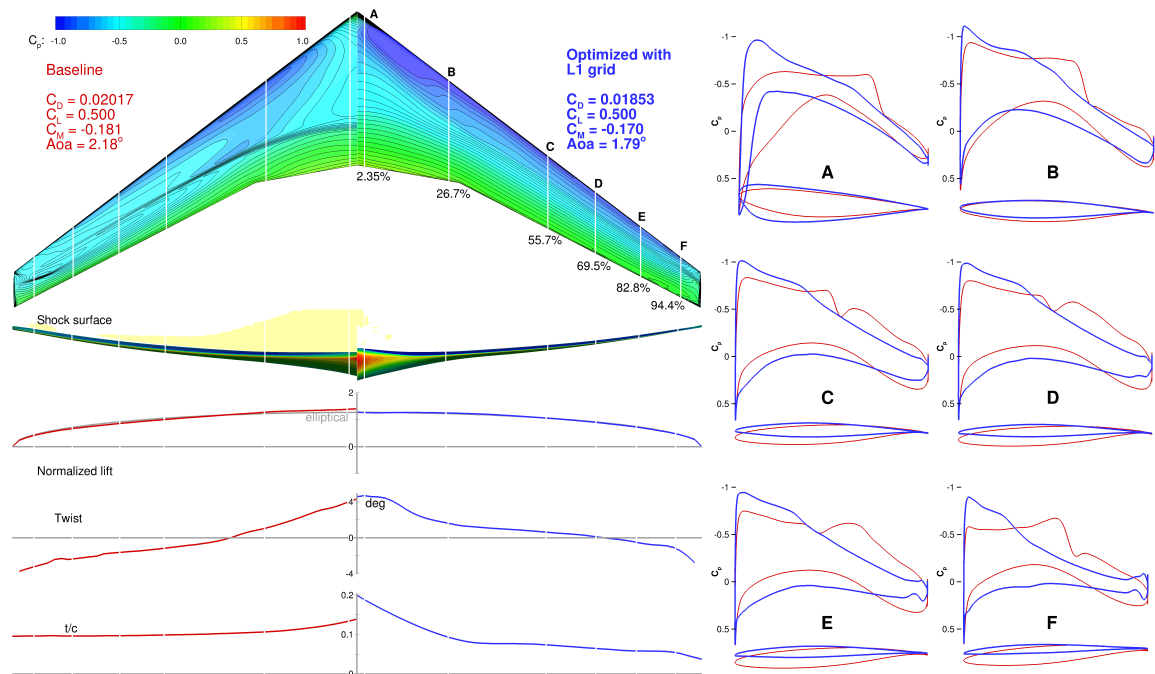


Figure 6.6: The optimized wing is shock-free and has 8.1% lower drag.

In this figure, the baseline wing results are shown in red and the optimized wing results are shown in blue. At the optimum, the lift coefficient target is met, and the pitching moment is reduced to the lowest allowed value. The lift distribution of the optimized wing is much closer to the elliptical distribution than that of the baseline, indicating induced drag is close to the theoretical minimum. This is achieved by fine-tuning the twist distribution and airfoil shapes. The baseline wing has a near-linear twist distribution. The optimized design has more twist at the root and tip, and less twist near mid-wing. The optimized thickness distribution is significantly different from that of the baseline, since the thicknesses are allowed to decrease to 25% of the original thickness, and there is a strong incentive to reduce the airfoil thicknesses in order to reduce wave drag. The volume is constrained to be greater than or equal to the baseline volume, so the optimizer drastically decreases the thickness of the airfoils on the outboard of the wing to the lower bounds, where there is less volume to be gained, while increasing the thickness near the root (up to 20%), where the chords are larger and the volume-drag trade-off is more favorable. The low outboard thickness would in practice incur a large structural weight penalty, and to trade off the reduction in drag and increase in weight would require aerostructural optimization [59].

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

This optimization uses a relatively large grid size (3.6M). It took a significant amount of computational cost: it converges after 914 optimization iterations, but it takes 616.5 hours (26 days) on 64 processors. Therefore, the computational cost is prohibitive for large grid size. The following two sections presented two new approaches in an effort to reduce the computational cost.

6.2.2 Aerodynamic Shape Optimization Using the Richardson’s Extrapolation

Richardson’s extrapolation is commonly used to determine the errors in the spatial discretization in CFD [120]. The idea is to use Richardson’s extrapolation to obtain the solution values at zero-grid spacing from a series of grids with coarser

discretization. We first determine the order of accuracy using three solutions with a constant grid refinement ratio r , as follows:

$$p = \ln \left(\frac{f_3 - f_2}{f_2 - f_1} \right) / \ln(r) \quad (6.1)$$

For a second-order CFD solver, p should be close to 2 when the grids are well resolved. We can then compute the zero-grid spacing values using,

$$f_{h=0} = f_1 + \frac{f_1 - f_2}{r^p - 1} \quad (6.2)$$

We attempted to extend this idea to aerodynamic shape optimization. Both flow solutions and the adjoint sensitivities are evaluated on the two coarser grids. We then use Richardson’s extrapolation to compute the zero-grid spacing objective function, constraints, and gradient values. Since the computational cost of the two coarser grids combined is still much lower than that of the fine grid, the total computational cost should be reduced. However, when we implemented this approach with the CRM wing optimization problem, the optimization was terminated due to numerical difficulties with only 5 iterations. The drag coefficient only had negligible changes. We found that both objective and constraints from the flow solution follows the Richardson’s extrapolation’s assumption: the solution is globally second-order, in addition to being locally second-order, and the solution functionals were computed using consistent second-order methods [120]. However, the gradients do not have a consistent order.

To further examine the gradients on different grid levels, we plot the shape gradients for each grid level, as shown in Figure 6.7.

The gradient can be affected by local phenomena such as shocks and separation, which causes the inconsistency between the grid levels. Thus, the zero-grid spacing gradients cannot be computed with Richardson’s extrapolation, and the optimization cannot improve the design with incorrect gradients. We notice that the gradients between each grid are close.

As an alternative, we perform the same direct optimization as Section 6.2.1 but with gradients computed on coarser L2 mesh in an effect to reduce the computational cost of the gradients. The optimization is able to converge further with the coarse gradients. It terminated with 54 iterations due to numerical difficulties. Once it terminated, we continue the optimization with the fine L1 grid gradients. Figure 6.8 show the optimized results.




Figure 6.7: The grid convergence for the shape gradients at different locations on the wing.

We observed a similar optimal shape as that in Section 6.2.1. The optimized drag is higher by 0.5 counts. Both lift and moment constraints are met at the optimum. The shock is also eliminated, as shown on the airfoil C_p distributions. The computational cost is reduced when compared to the direct L1 optimization. The total time is 441.8 hours (18.4 days), resulting in a 28% reduction.

6.2.3 Aerodynamic Shape Optimization Using Multilevel Technique

In this section, we present an acceleration technique inspired by the multigrid start-up procedure in CFD that reduced the overall computational cost of the optimization effectively. Since we have improved the efficiency of our flow and adjoint solvers significantly over the last few years [18, 28, 58], we seek new methods to further reduce the computational cost of the aerodynamic shape optimization.

Since it is less costly to compute both the flow solution and the gradient on a coarser grid, we perform the optimization first on the coarsest grid until a certain

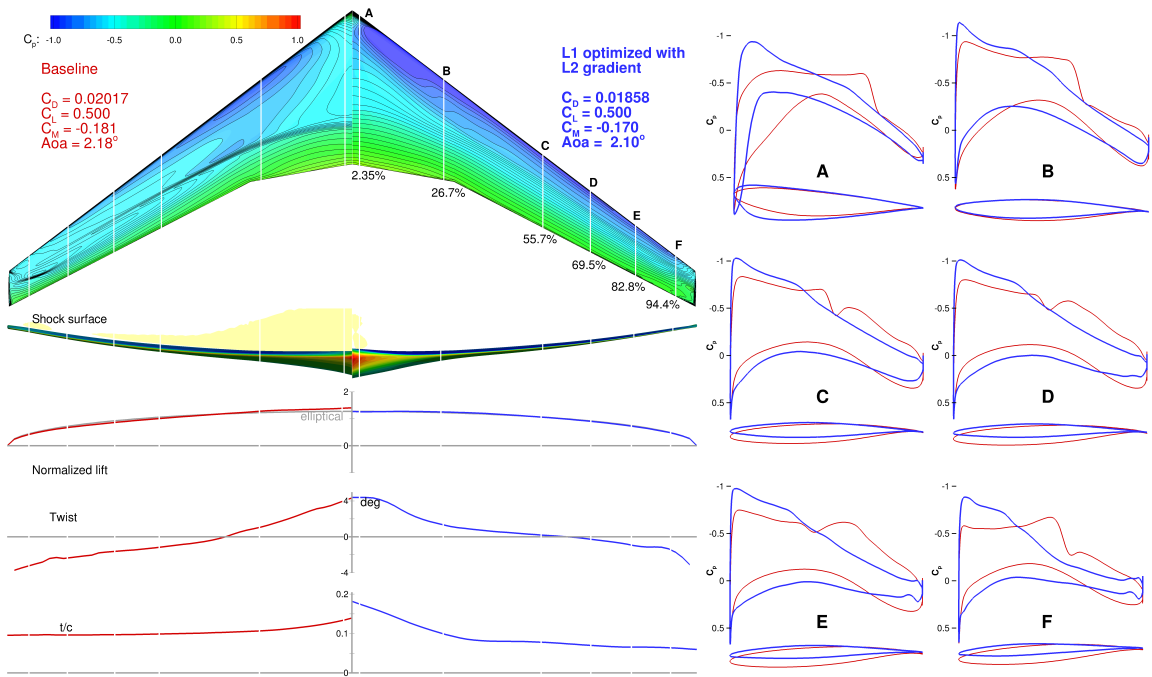


Figure 6.8: The optimization starts with the coarser L2 gradients.

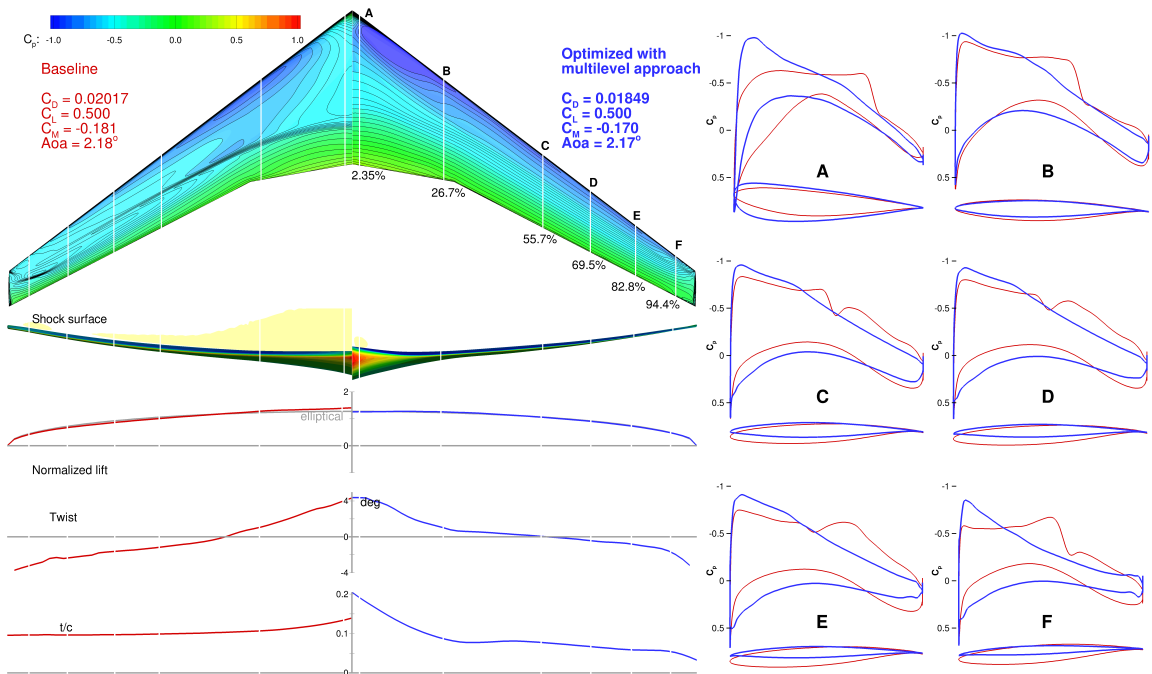


Figure 6.9: Multilevel optimization with L2 and L1 grids significantly reduced computational cost.

level of optimality is achieved. Then, we move to the next grid level and start with the optimal design variables from the coarser grid. Since the drag and lift coefficients are generally different for each grid level, the approximate Hessian (used by the gradient-based optimizer) must be restarted. We repeat this process until the optimization on the finest grid has converged. Note that this procedure is different from traditional multigrid methods, where the coarse levels are revisited via multigrid cycles.

We performed the same CRM wing optimization problem using the approach described above. Figure 6.9 shows the comparison of the baseline and the optimized results. The optimized design is very similar to that in the previous two sections. The difference between the drag coefficients is within one count. There are visible differences in the airfoils C_p distributions, as shown in Figure 6.10. This might be caused by local minima that are close to each other, as previously observed by Lyu *et al.* [115, 101]. The multilevel approach has also been successfully applied in [101].

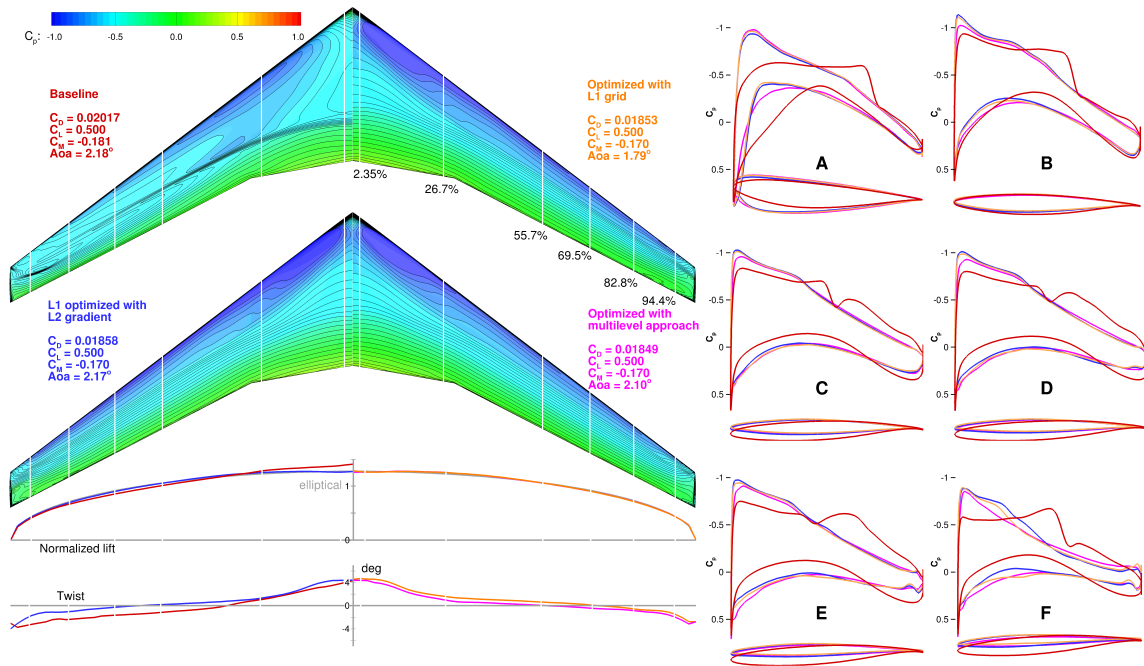


Figure 6.10: The difference in drag coefficients is within 1 count.

The multilevel approach uses significantly less computational resources. We use two grid levels: L2 (451 k cells) and L1 (3.6 M cells). We can see that most of the optimization iterations are performed on the coarse grid, and as a result, the number of the function and gradient evaluations on the successively finer grids is greatly reduced. The optimization has 638 iterations on the L2 grid. Thanks to the optimization

with the coarser grid, only 89 iterations are needed on the L1 grid to converge the optimization. The optimization with multilevel approach converges in 95.4 hours (4 days) on 64 processors. Table 6.3 shows the comparison of the optimization time. The merit function, optimality, and feasibility histories are plotted in Figure 6.11; detailed definitions of these values can be found in the SNOPT manual [121].

	Iterations	Time (hr)	Reduction
L1 direct	914	616.5	
L1 with L2 gradients	645	441.8	-28.3%
L1 multilevel	727	95.4	-84.5%

Table 6.3: The optimization with the multilevel approach reduces computational time by 84.5%.

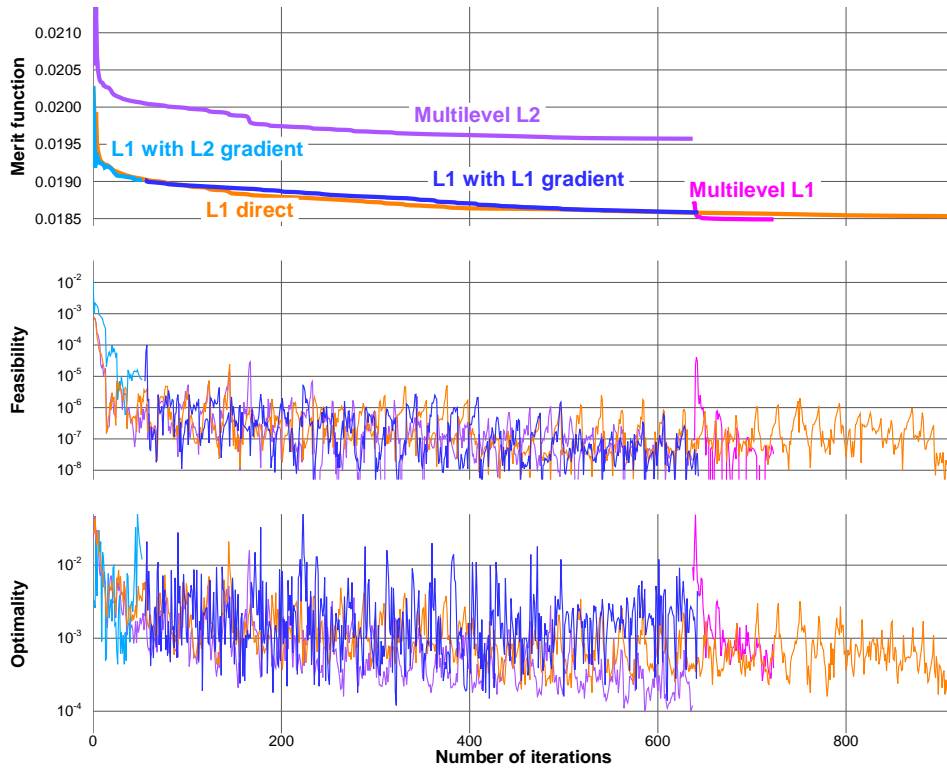


Figure 6.11: The optimization history of the three approaches.

We can see that the multilevel approach achieves the lowest computational cost by a large margin. With the increase of the grid size, this benefit becomes even larger. In this approach, we only performed one-way grid sequencing. Additional benefits may be achieved with a true multigrid V- or W-cycle at the optimization level.

6.3 Conclusions

In this chapter, we have presented a study of the CRM wing shape optimization problem using three different optimization approaches. The optimization problem is defined by the Aerodynamic Design Optimization Discussion Group (ADODG). The drag coefficient is minimized for one flight condition with respect to 720 shape design variables, subject to lift, pitching moment, and geometric constraints. We compared direct optimization with optimization using Richardson's extrapolation, and with optimization using multilevel approach. We found that the multilevel approach achieved the lowest computational cost. The total computation time was reduced from 616.5 hours to 95.4 hours using two grid levels.

The strategies presented in this chapter open a new door to aerodynamic shape optimization. Further development of the techniques at the optimization level, in conjunction with MDO architectures have the potential to make future large-scale optimization more efficient and effective.

CHAPTER 7

Aerodynamic Design Optimization of a Current-Generation Aircraft

Despite considerable research on aerodynamic shape optimization, there is no standard benchmark problem allowing researchers to compare results. This chapter addresses this issue by solving a series of aerodynamic shape optimization problems based on the Common Research Model wing benchmark case defined by the Aerodynamic Design Optimization Discussion Group (ADODG) using the aerodynamic shape optimization framework developed in this thesis. The aerodynamic model solves the Reynolds-averaged Navier–Stokes equations with a Spalart–Allmaras turbulence model. A gradient-based optimization algorithm is used in conjunction with an adjoint method that computes the required derivatives.

The majority of the aerodynamic shape optimization problems in the literature are solved with gradient-based optimizers [3, 4, 11, 14, 100]. High-fidelity aerodynamic shape optimization with large number of design variables has the potential to have multiple local minima. The problem is that due to the high number of dimensions and the high cost of the function evaluations, the design space is impossible to visualize fully. This makes it challenging estimate the number of local minima and to form a complete picture of the design space. Several authors explored the multimodality in aerodynamic shape optimization with gradient-free optimization [122, 103], and combinations of gradient-free and gradient-based optimization [123]. However, there has been no thorough study for RANS-based three-dimensional aerodynamic shape optimization with large numbers of shape variables. In this chapter, we also explore multimodality by performing several shape optimizations starting from randomly generated geometries. This work is based on previous papers presented by the author [115, 101].

7.1 Optimization Problem Formulation

The optimization problem formulation is the same as that presented in the previous chapter. The aerodynamic shape optimization seeks to minimize the drag coefficient by varying the shape design variables subject to a lift constraint ($C_L = 0.5$) and a pitching moment constraint ($C_{My} \geq -0.17$). Shape design variables are the z -coordinate movements of 720 control points on the FFD volume (shown in Fig. 6.2) and the angle-of-attack. Control points at the trailing edge are constrained to avoid any movement of the trailing edge. Therefore, the twist about the trailing edge can be implicitly altered by the optimizer using the remaining degrees of freedom. The leading-edge control points at the wing root are also constrained to maintain a constant incidence for the root section. There are 750 thickness constraints imposed in a 25 chordwise and 30 spanwise grid covering the full span and from 1% to 99% local chord. The thickness is set to be greater than 25% of the baseline thickness at each location. Finally, the internal volume is constrained to be greater than or equal to the baseline volume. The complete optimization problem is described in Table 6.2.

7.2 Single-Point Aerodynamic Shape Optimization

In this section, we present our aerodynamic design optimization results for the CRM wing benchmark problem (described in Table 6.2) under the nominal flight condition (Mach 0.85, $Re = 5 \times 10^6$). We use the L0 grid (28.8M cells) for the optimization, thanks to a multilevel optimization acceleration technique that significantly reduces the overall computational cost of the optimization. The details of this technique are presented in Sec. 7.3. Our optimization procedure reduced the drag from 199.7 counts to 182.8 counts, i.e., an 8.5% reduction. The corresponding Richardson-extrapolated zero-grid spacing drag decreased from 199.0 counts to 181.9 counts. Given that the CRM configuration was designed by experienced aerodynamicists, this is a significant improvement (although they designed the wing in the presence of the fuselage, which we are ignoring in this problem).

Figure 7.1 shows a detailed comparison of the baseline wing and the optimized wing. In this figure, the baseline wing results are shown in red and the optimized wing results are shown in blue. At the optimum, the lift coefficient target is met, and the pitching moment is reduced to the lowest allowed value. The lift distribution of the optimized wing is much closer to the elliptical distribution than that of the baseline, indicating an induced drag that is close to the theoretical minimum for a

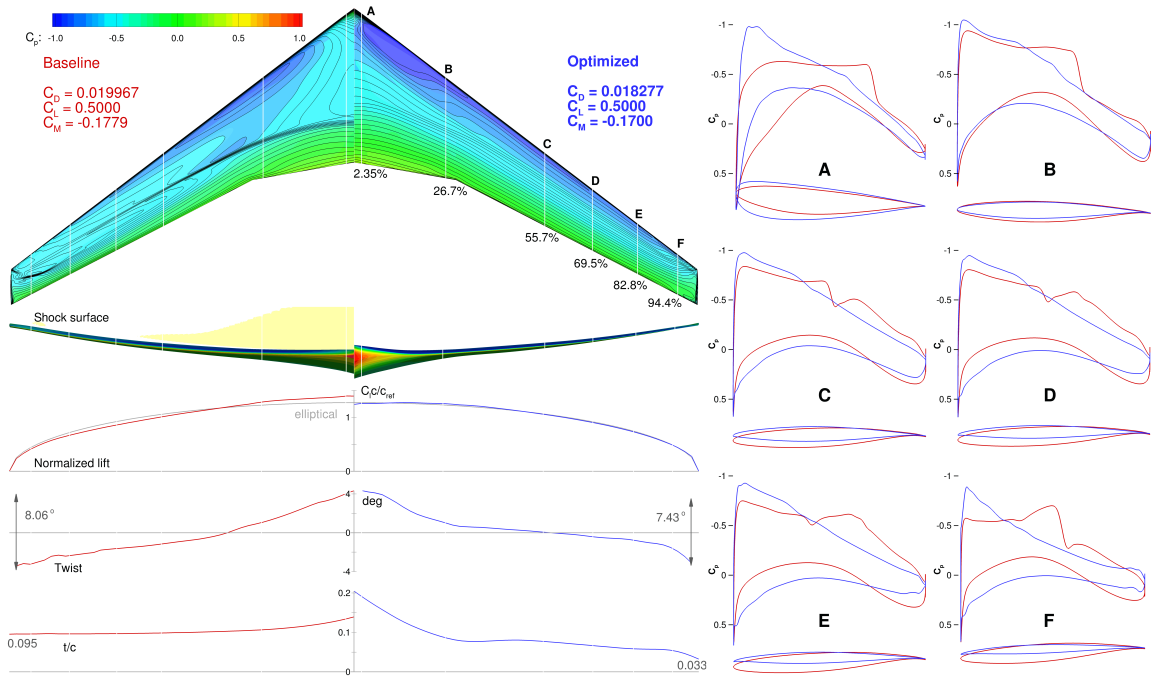


Figure 7.1: The optimized wing is shock-free and has 8.5% lower drag.

planar wake. This is achieved by fine-tuning the twist distribution and airfoil shapes. The baseline wing has a near-linear twist distribution. The optimized design has more twist at the root and tip, and less twist near mid-wing. The overall twist angle changed only slightly: from 8.06 to 7.43 degree. More detailed breakdowns of drag components can be obtained using [124].

The optimized thickness distribution is significantly different from that of the baseline, since the thicknesses are allowed to decrease to 25% of the original thickness, and there is a strong incentive to reduce the airfoil thicknesses in order to reduce wave drag. Volume is constrained to be greater than or equal to the baseline volume, so the optimizer drastically decreases the thickness of the airfoils on the outboard of the wing to the lower bounds, where there is less volume to be gained, while increasing the thickness near the root (up to 20%), where the chords are larger and the volume-drag trade-off is more favorable. Telidetzki *et al.* [117] observed similar trends in their results. The low outboard thickness would in practice incur a large structural weight penalty, and to trade off the reduction in drag and increase in weight would require aerostructural optimization [59]. To obtain a more realistic design without resorting to aerostructural optimization, in Sec. 7.6 we solve an additional optimization problem with a stricter thickness constraint.

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

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

Another noticeable feature in the optimized wing is the sharp leading edges in the outboard wing sections. The optimizer exploits a weakness in the problem formulation: with a single-point optimization, there is no penalty for thinning out the leading edge. In practice, however, sharp-leading-edge airfoils experience adverse performance under off-design conditions, since the flow is prone to separation at off-design angles-of-attack. We address these issues in more detail by performing a multipoint optimization in Sec. 7.7.

To ensure that the result of our single-point optimization has sufficient accuracy, we conducted a grid convergence study of the optimized design. Table 7.1 summarizes the results for each grid level. The mesh convergence plot for both the baseline and optimized geometry meshes is shown in Fig. 6.3. The zero-grid spacing drag, which was obtained using Richardson’s extrapolation, is also plotted in the figure. We can see that the L0 mesh has sufficient accuracy: the difference in the drag coefficient for the L0 mesh and the value obtained for the zero-grid spacing is within one drag count. The variation in drag coefficient between the baseline and optimized meshes is nearly constant for each grid level, which gives us confidence that the optimization using the coarse meshes represent the design space trends sufficiently well. Therefore, we perform the remaining optimization studies on the coarser mesh (L2), assuming that we capture the correct design trends.

Mesh level	Mesh size	Baseline C_D	Optimized C_D	ΔC_D	C_L	Optimized C_M	Optimized α
$h = 0$	∞	0.01990	0.01819	0.00171			
L00	230, 686, 720	0.01992	0.01820	0.00171	0.5000	-0.1694	2.1759 degree
L0	28, 835, 840	0.01997	0.01825	0.00172	0.5000	-0.1700	2.1660 degree
L1	3, 604, 480	0.02017	0.01846	0.00171	0.5000	-0.1710	2.1584 degree
L2	450, 560	0.02111	0.01964	0.00147	0.5000	-0.1731	2.1970 degree

Table 7.1: The drag differences between the baseline and optimized meshes are nearly constant for each grid level.

7.3 Multilevel Optimization Acceleration Technique

To speed up the convergence of this optimization, we use the multilevel acceleration technique that presented in the previous chapter. We used this procedure to obtain the optimal wing presented in the previous section. We use three grid levels: L2 (451 k cells), L1 (3.6 M cells), and L0 (28.8 M cells). The merit function, optimality, and feasibility histories are plotted in Fig. 7.2; detailed definitions of these values can be found in the SNOPT manual [121]. We can see that most of the optimization iterations are performed on the coarse grid, and as a result, the number of the function and gradient evaluations on the successively finer grids is greatly reduced. Table 7.2 summarizes the computational time spent on each grid level. Thanks to the optimization with the coarser grids, only 18 iterations are needed on the L0 grid to converge the optimization. However, the L0 grid requires the largest computational effort, due to the high cost of the flow and adjoint solutions on this fine grid. Given that the cost per optimization iteration in the L0 grid is 770 proc-hr (compared to 2.9 proc-hr for the L2 grid) it is not feasible to perform an optimization using only the L0 grid. Assuming that the same number of iterations used for the L2 grid (638) would be needed for the L0 grid, the computational cost would be 23 times higher than that of the multilevel approach, which would correspond to 16 days using 1248 processors.

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

Table 7.2: The number of iterations on the L0 grid is reduced to 18.

Figure 7.3 shows the initial and optimized results at each grid level. If we examine the results more closely, we see that the optimized results for the L2, L1, and L0 grids are all similar. This validates the underlying assumption of this method: that a coarser grid provides a good approximation to the design space of the finer grid when the set of design variables remains the same. Most of the computational effort on the subsequent grid levels is spent on smoothing out the shock that reappeared because of the finer grid spacing. This multilevel acceleration technique proved to significantly reduce the number of iterations needed to optimize in the fine grid, and the total computational effort was greatly reduced.




Figure 7.2: Most of the computations are performed on the coarse grid.

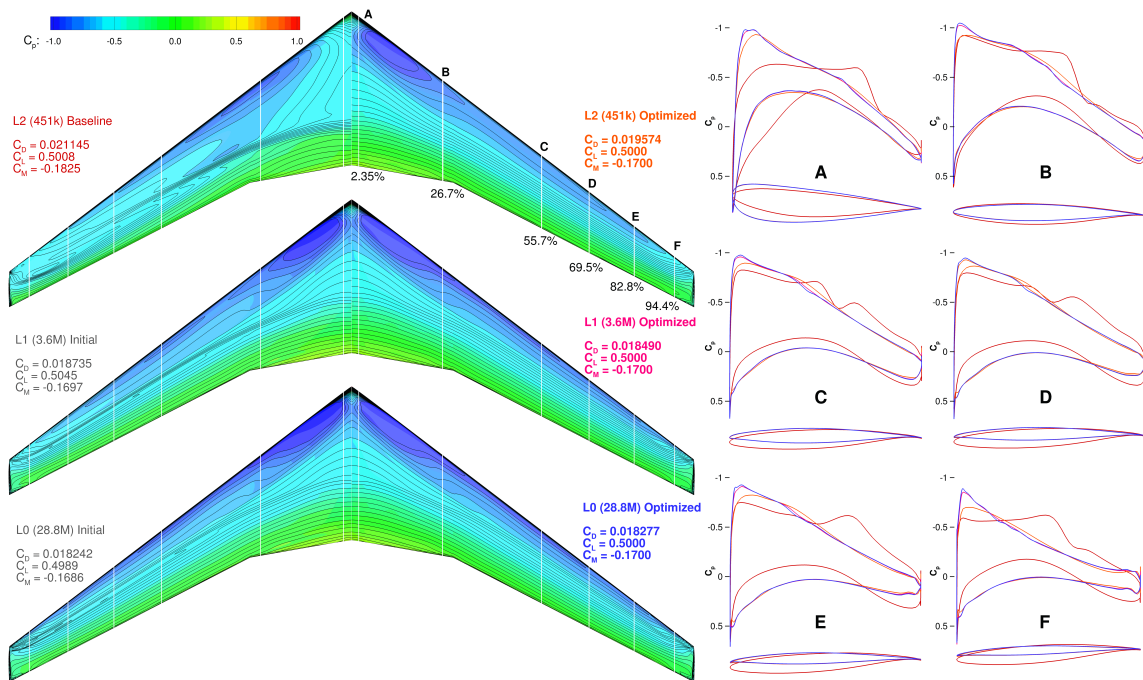


Figure 7.3: The optimized results of each grid level exhibit only subtle differences.

7.4 Aerodynamic Shape Optimization Starting from a Random Geometry

The existence of multiple local minima in RANS-based three-dimensional aerodynamic shape optimization with respect to large numbers of design variables has yet to be explored. The problem is that due to the high number of dimensions, the design space is difficult to visualize. In addition, the function evaluations are costly, making it challenging to explore the design space thoroughly and come to definitive conclusions.

We explore the multi-modality of the single-point aerodynamic shape optimization problem described in Section 7.2, by solving separate optimizations starting from four different geometries. The first starting geometry is the CRM wing of Sec. 7.2. The other three starting geometries are randomly generated by applying a random perturbation to each design variable of the CRM wing, resulting in completely different geometries. The volume constraint is imposed, such that the volume of the baseline CRM wing is preserved. The initial starting points for the three random runs are shown in Fig. 7.4. C_p distribution is shown on the surface, along with a visualization of the shock (orange) and separation (red).

Figure 7.5 shows the optimized results from a random initial geometry. The optimization is performed on the L2 grid. We can see that the performance of the initial design is extremely poor. This is no surprise, since the airfoil shapes are unlike anything one would design: they exhibit oscillations and sharp edges, resulting in a wildly varying C_p distribution. In addition, the flow solution is probably not accurate. In spite of these wild shapes and the inaccuracy of the flow solution, the gradients seem to point in the right direction, since the optimizer is able to smooth out the airfoils and achieve a shock-free wing similar to the original single-point design presented in Sec. 7.2. All the constraints are met, and the lift distribution is close to elliptical. This optimization demonstrates the robustness of our aerodynamic optimization approach and showcases the power of the adjoint method.

We performed the same optimization for three random starting points and compared the results against each other, as well as against the single-point optimized wing, as shown in Fig. 7.6. Each optimized result is color-coded, and the nominal optimized result from Sec. 7.2 is shown in black. Overall, there are only small differences between the four designs, as evidenced by the similar C_p distributions and cross sectional shapes. The difference in drag between all four designs is within one drag count. However, there are still some small visible differences, indicating the possibility that the design space is multimodal.

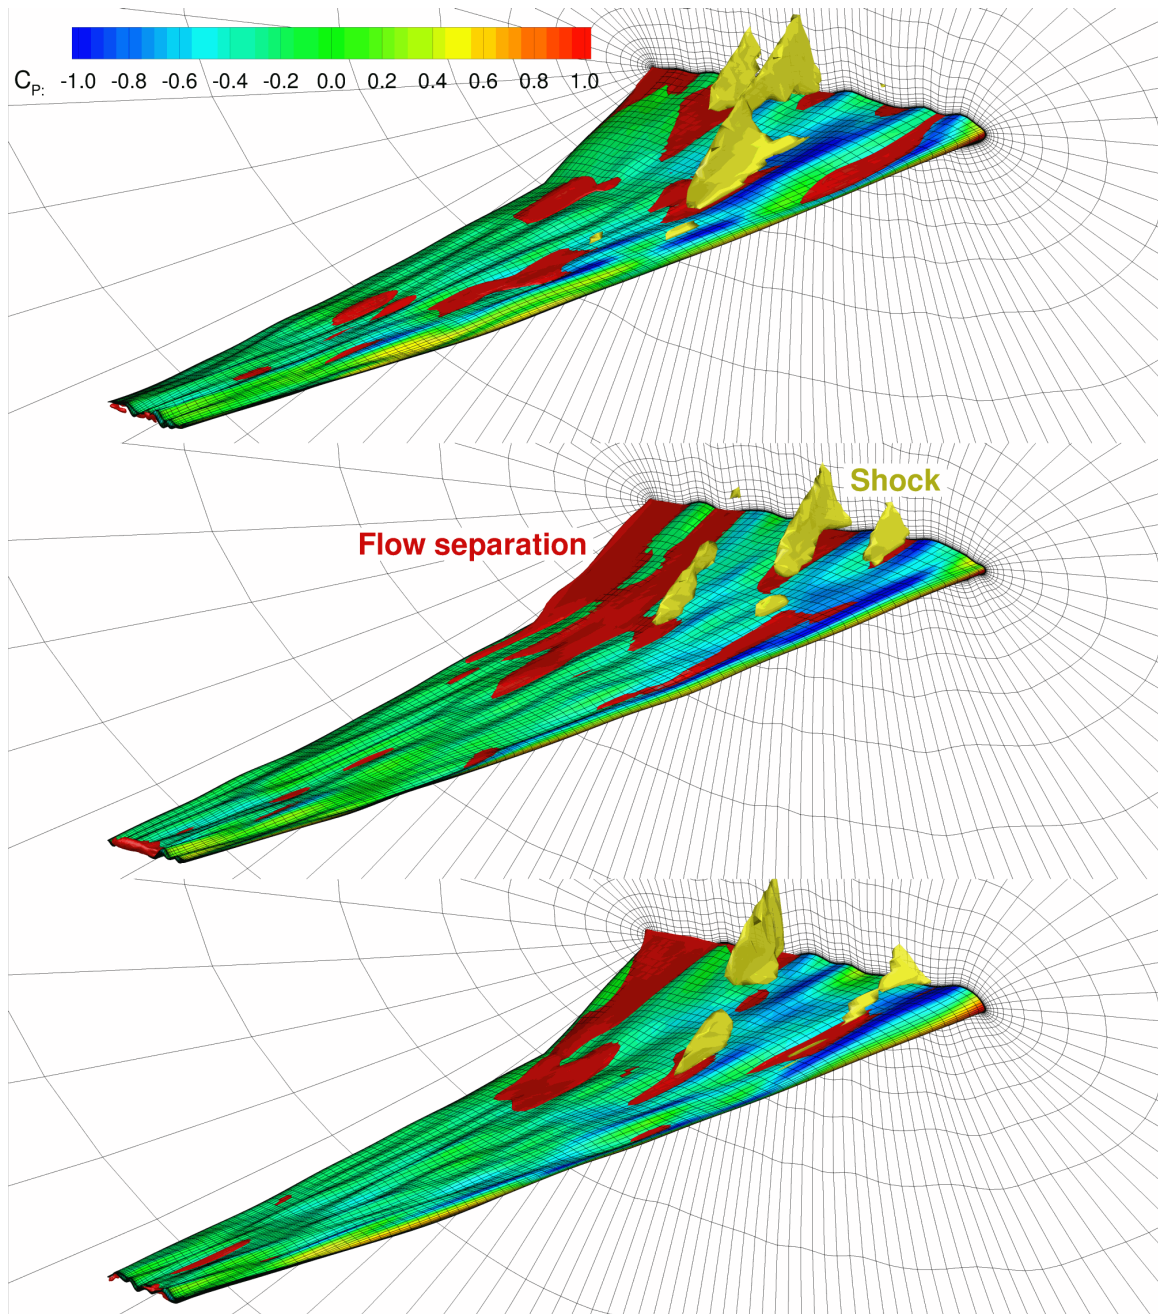


Figure 7.4: The initial geometries are randomly generated from the baseline CRM wing.

To further visualize this design space, we compute the merit function in the design space between two optimized designs, as shown in Figure 7.7. The merit function is a combination of the objective function and the constraints. [121] We are able to visualize a slice of the design space by plotting the merit function along a line between two optima. A series of wing shapes are generated by linearly varying all of the design

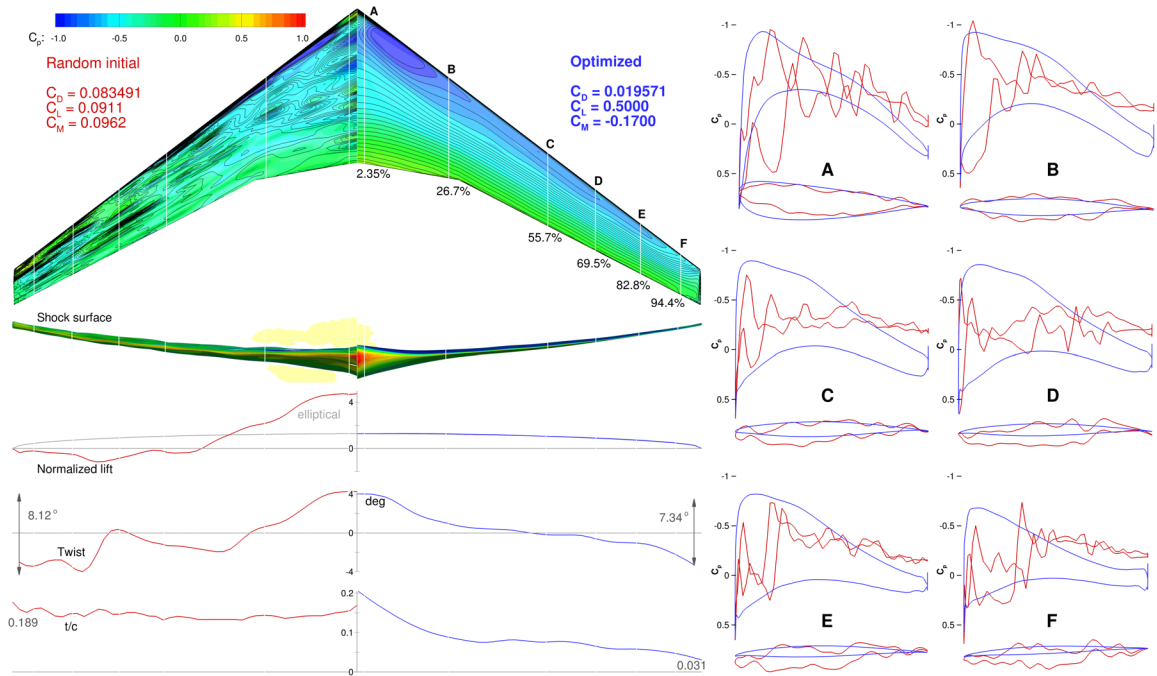


Figure 7.5: The optimization manages to start from a random geometry and converge to an optimal wing that is shock free.

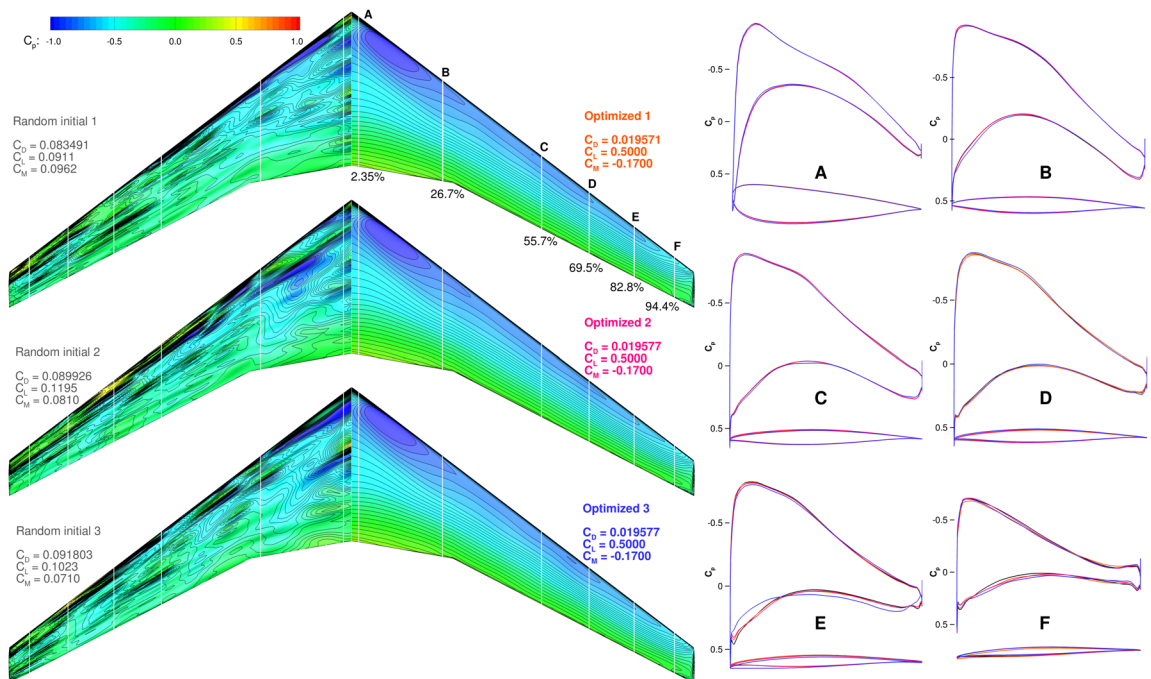


Figure 7.6: All three optimizations with random starting geometries converged to similar optima.




Figure 7.7: The merit function values between optimized designs show the local minima.




Figure 7.8: The Euclidean distances between the multiple local minima are similar and are all under 3.4 in (1.2% of the mean aerodynamic chord).

variables. A CFD solution is solved for each of those designs to obtain the merit function. As shown in this figure, the merit function does appear to have multiple local minima, even though the values of the merit function are within one count among those optima. In addition, we also computed the mean difference between the design variables of each of the optimized designs as shown in Fig. 7.8. The three optima appear to be nearly equally spaced in the design space, with a Euclidean distance ranging from 2.72 to 3.34in, which corresponds to only 1.2% of the mean aerodynamic chord. Based on this data, we believe that the design space for this aerodynamic shape optimization problem is mostly convex, but that it has a small flat region that is multimodal. The humps and local minima could also be caused by the constraints.

7.5 Effect of the Number of Shape Design Variables

The cost of computing gradients with an efficient adjoint implementation is nearly independent of the number of design variables. We took advantage of this efficiency by optimizing with respect to 720 shape design variables in the previous sections. However, we would like to determine the trade-off between the number of design variables and the optimal drag, and to examine the effect on the computational cost of the optimization. Thus, in this section we examine the effect of reducing the number of design variables. A series of new enlarged FFDs are created to ensure proper geometry embedding for small numbers of design variables. The shape design variables are distributed in a regular grid, where the finest grid has $15 \times 48 = 720$ design variables. The 15 chordwise stations correspond to 15 distinct airfoil shapes, while the shape of each airfoil is defined by 48 control points (half of these on the top, and the other half on the bottom).

We solve the optimization problem of Sec. 7.2 using the L2 grid with variations in the number of defining airfoils and the number of points per airfoil. Figure 7.9 shows the resulting optimized designs for different numbers of airfoil control points and a fixed number of defining airfoils. Reducing the airfoil control points from 48 to 24 has a negligible effect on the optimal shape and pressure distribution, and the optimum drag increases by only 0.1 counts. As we further reduce the number of airfoil points to 12 and 6, both the drag and pressure distribution show noticeable differences.

Variation in the number of defining airfoils follows a similar trend to the variation in the number of airfoil control points, as shown in Fig. 7.10. However, the drag penalty due to the number of airfoils is less severe than the penalty observed in the

airfoil point reduction. Therefore, increasing the number of design variables in the chordwise direction is more beneficial than increasing the number of defining airfoils in the spanwise direction.

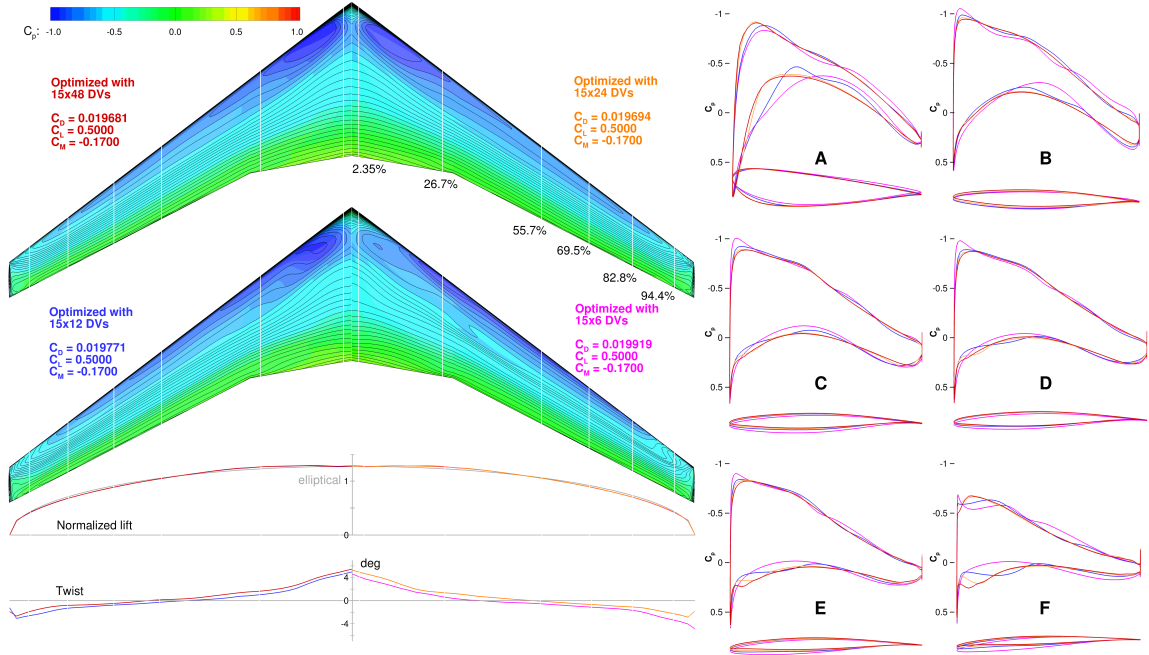


Figure 7.9: Optimized designs with varying number of airfoil control points.

We also perform the optimization with a reduced number of shape design variables in both the chordwise and spanwise directions simultaneously, as shown in Fig. 7.11. From this study we conclude that an adequate optimized design can be achieved with a smaller number of design variables: with $8 \times 24 = 192$ shape variables, the difference in the optimal drag coefficient is only 0.6 counts. Any further reduction in the number of design variables has a much larger impact on the optimal drag.

Figure 7.12 plots the convergence history for each optimization case. When we decrease the number of airfoil control points, the number of optimization iterations required decreases drastically. However, the number of defining airfoils has little effect on the optimization effort. This is in part because the adjoint computational cost is independent of the number of design variables. In addition, the coupled effects between design variables are much stronger between variables within an airfoil than between variables in different airfoils.

For an optimization process in which the computational cost scales with the number of design variables, such as when the gradients are computed via finite differences, or for gradient-free optimizers, a smaller number of design variables would signifi-

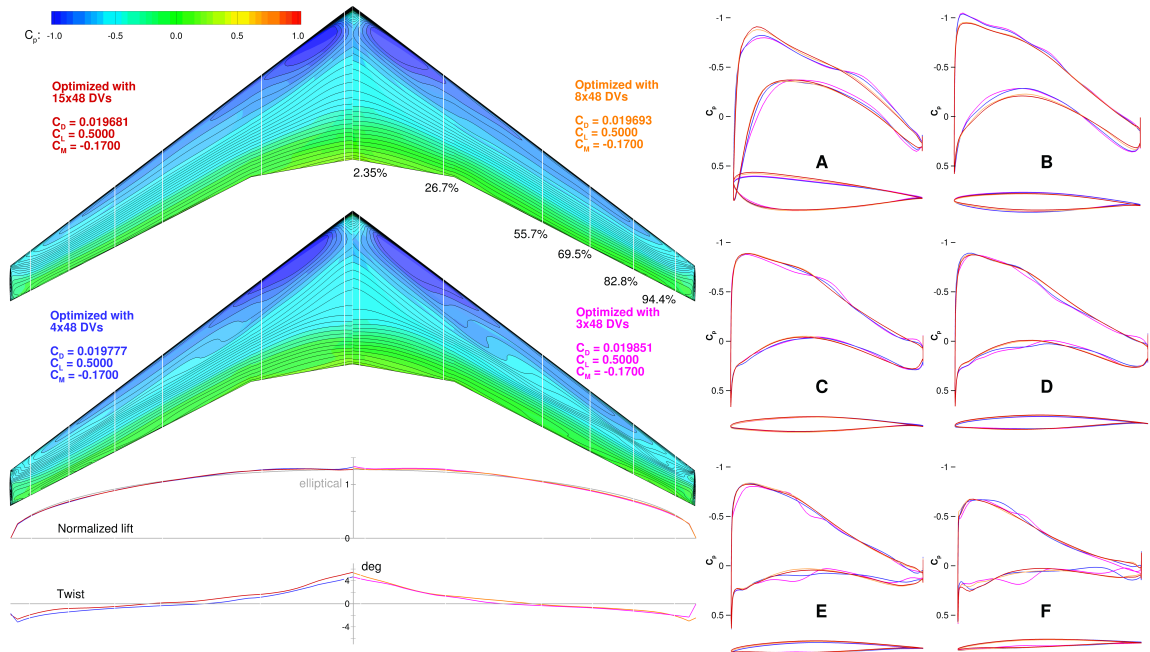


Figure 7.10: Optimized designs with varying number of airfoil sections.

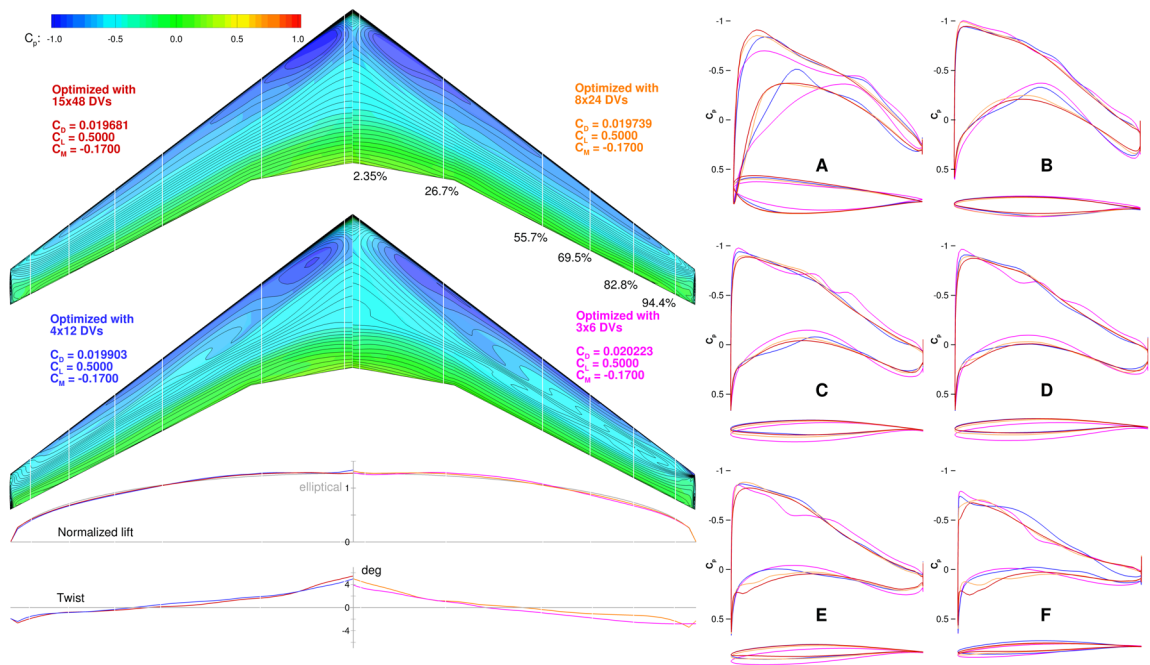


Figure 7.11: Optimized designs for varying numbers of shape design variables.

cantly impact the optimized design. For example, for $3 \times 6 = 18$ variables, the drag of the optimized design would increase by 5.4 counts.

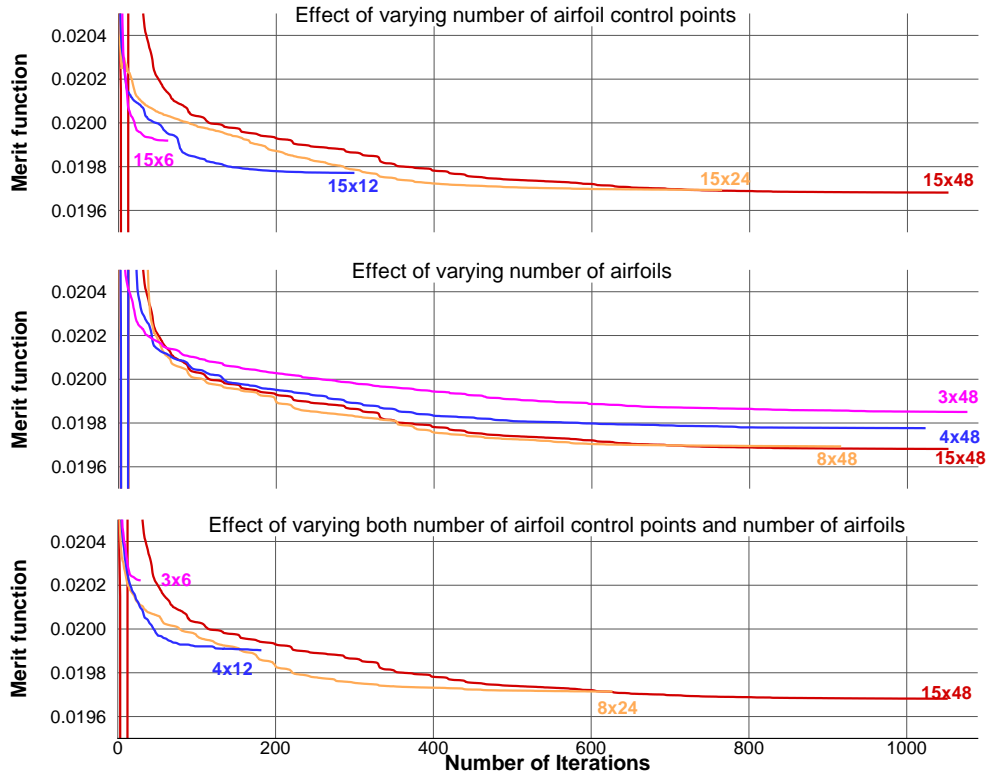


Figure 7.12: The number of optimization iterations does not decrease significantly as the number of defining airfoils is decreased.

7.6 Aerodynamic Shape Optimization without Thickness Reduction

As seen in Sec. 7.2, the optimized wing has a thickened root airfoil and an unrealistically thin tip airfoil. To address this issue, we solved an optimization problem identical to that solved in Sec. 7.2 except for modified thickness constraints: all thicknesses must be greater than or equal to the baseline thickness (instead of being allowed to decrease to 25% of the baseline thickness). The optimization is performed on the L2 grid, and the results are shown in Fig. 7.13.

The results of the optimization with no thickness reduction are shown in black. The spanwise lift and twist distributions for the two cases are similar. However, the pressure distribution and airfoil shapes are significantly different, especially those near the wing root and wing tip. The mean difference between the baseline and optimized designs is only 1.1 inches. The optimized wing with no thickness reduction has five additional drag counts when compared with the optimized wing that allowed 25%




Figure 7.13: The drag on the optimized wing is five counts higher if no airfoil thickness reduction is allowed.

of the baseline thickness. This aerodynamic performance penalty may be compensated for by the reduction in the wing weight when structural design is considered. A detailed aerostructural optimization would be necessary to examine the multidisciplinary trade-offs involved [58, 59]. In addition, the optimization takes significantly fewer iterations (296 iterations) as compared to the optimization in Sec. 7.2 (638 iterations). This is due to the absence of volume-thickness trade resulting from the tighter thickness constraints.

7.7 Multipoint Aerodynamic Shape Optimization

Transport aircraft operate at multiple cruise conditions because of variability in the flight missions and air traffic control restrictions. Single-point optimization under the nominal cruise condition could overstate the benefit of the optimization, since the optimization improves the on-design performance to the detriment of the off-design performance. In Sec. 7.2, the single-point optimized wing exhibited an unrealistically sharp leading edge in the outboard of the wing. This was caused by a combination of the low value for the thickness constraints (25% of the baseline) and the single-point formulation.

A sharp leading edge is undesirable because it is prone to flow separation under off-design conditions. We address this issue by performing a multipoint optimization. The optimization is performed on the L2 grid. We choose five equally weighted flight conditions with different combinations of lift coefficient and the Mach number, as previously done by the authors [59]. The flight conditions are the nominal cruise, $\pm 10\%$ of cruise C_L , and ± 0.01 of cruise Mach, as shown in Fig. 7.14. More sophisticated ways of choosing multipoint flight conditions and their associated weights can be used, such as the automated procedure developed by Liem *et al.* [125] that minimizes fleet-level fuel burn. The objective function is the average drag coefficient for the five flight conditions, and the moment constraint is enforced only for the nominal flight condition.

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

Table 7.3: The multiple flight conditions represent a five-point stencil in Mach- C_L space.




Figure 7.14: The multipoint optimization flight conditions represent a five-point stencil in Mach- C_L space.

A comparison of the single-point and multipoint optimized designs is shown in Fig. 7.15. The single-point results are shown in blue, and the multipoint results are

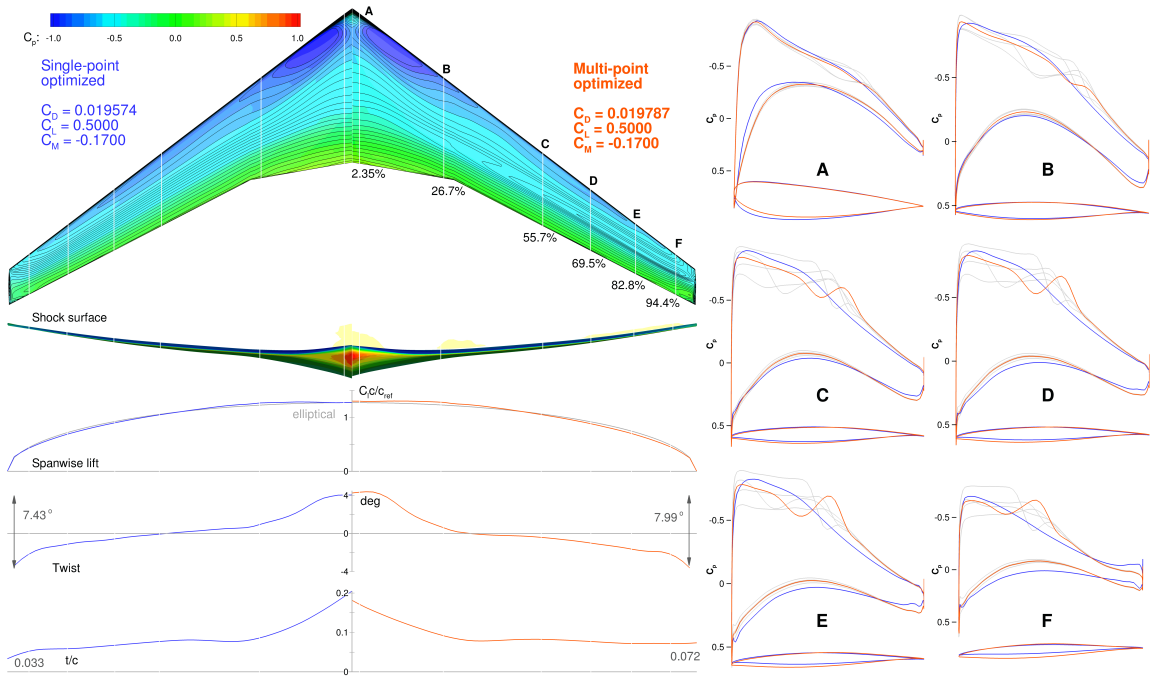


Figure 7.15: The multipoint optimized wing has a weak shock on the upper surface for each flight condition.

shown in orange. The C_p for the multipoint optimized result corresponds to the nominal condition. The multipoint sectional C_p of flight conditions 2–5 are plotted in gray. Unlike the shock-free design obtained with single-point optimization, the multipoint optimization settled on an optimal compromise between the flight conditions, resulting in a weak shock at all conditions. The leading edge is less sharp than that of the single-point optimized wing. Additional flight conditions, such as a low-speed flight condition, would be needed to further improve the leading edge. The overall pressure distribution of the multipoint design is similar to that of the single-point design. The twist and lift distributions are nearly identical. Most of the differences are in the chordwise C_p distributions in the outer wing section. The drag coefficient under the nominal condition is approximately two counts higher. However, the performance under the off-design conditions is significantly improved. Similar trends were observed in the multipoint optimization of Vassberg *et al.* [116].

To demonstrate the robustness of the multipoint design, we plot ML/D contours of the baseline, single-point, and multipoint designs with respect to C_L and cruise Mach in Fig. 7.16. ML/D provides a metric for quantifying aircraft range based on the Breguet range equation with constant thrust-specific fuel consumption. While the thrust-specific fuel consumption is actually not constant, assuming it to be constant

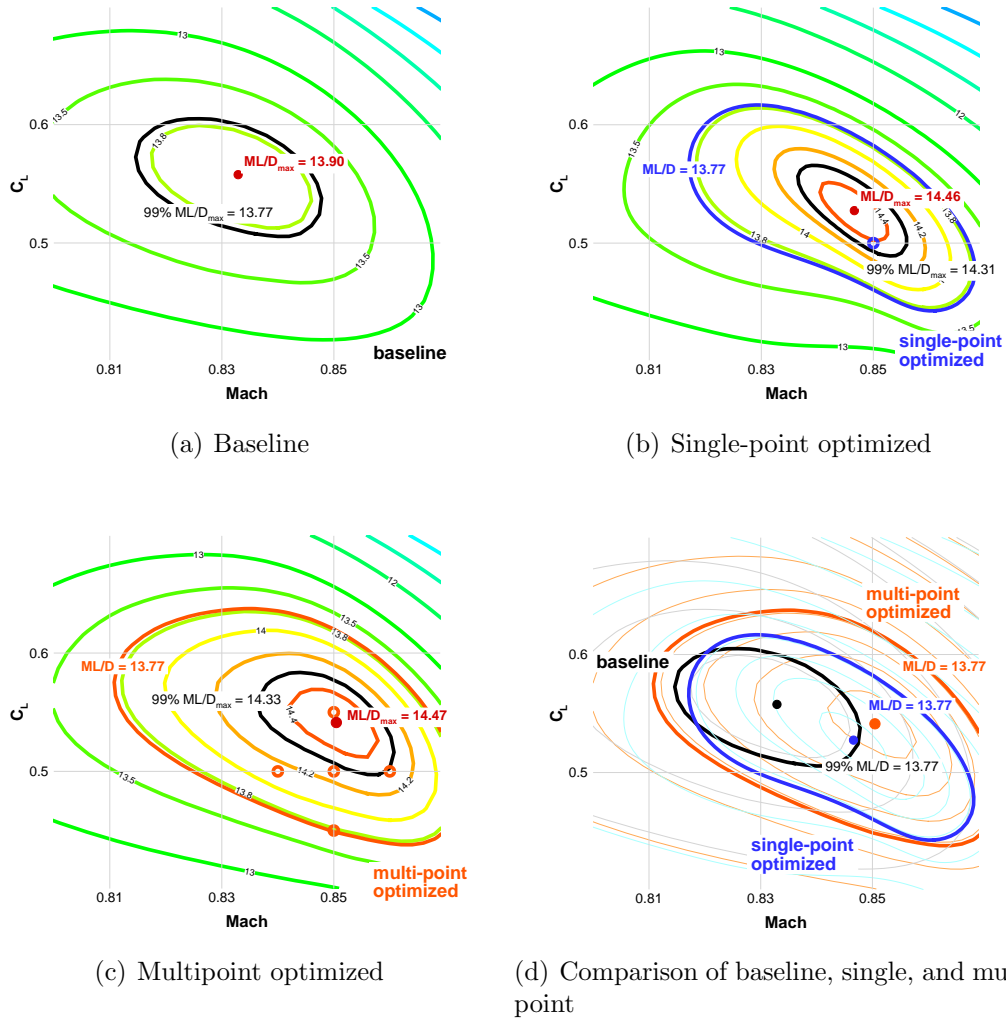


Figure 7.16: The multipoint optimized wing has better off-design performance and is more robust to changes in flight conditions.

is acceptable when comparing performance in a limited Mach number range [126]. We add 100 drag counts to the computed drag to account for the drag due to the fuselage, tail, and nacelles, and we get more realistic ML/D values.

The baseline maximum ML/D is at a lower Mach number and a higher C_L than that of the nominal flight condition. The single-point optimization increases the maximum ML/D by 4% and moves this maximum toward the nominal cruise condition. If we examine the variation of ML/D along the $C_L = 0.5$ line, we see that the maximum occurs at the nominal Mach of 0.85, which corresponds to a dip in a drag divergence plot.

For the multipoint optimization, the optimized flight conditions are distributed

in the Mach- C_L space, resulting in a flattened ML/D variation near the maximum, which means that we have more uniform performance for a range of flight conditions. In aircraft design, the 99% value of the maximum ML/D contour is often used to examine the robustness of the design [37]. The point with the highest Mach number on that contour line corresponds to the long range cruise (LRC) point, which is the point at which the aircraft can fly at a higher speed by incurring a 1% increase in fuel burn [127]. In this case, we see that the 99% value of the maximum ML/D contour of the multipoint design is larger than that of the single-point optimum, indicating a more robust design.

The ML/D_{\max} of the multipoint design is slightly higher than the maximum for the single-point design. While this seems counter-intuitive, it can be explained by the fact that the analysis conditions do not line up with the optimized maximum location. The optimizer has no information about the exact flight condition of the peak location, and thus it does not directly control the value at that location. A potential remedy for this mismatch would be the addition of two degrees of freedom to the optimization problem: the nominal Mach number and the nominal lift coefficient. This would allow the optimizer to track the ML/D_{\max} location during the optimization; upon convergence, the objective value would reflect the maximum possible performance in the $M-C_L$ space. Performing such an optimization, however, would require propulsion and operating cost models, and would involve multidisciplinary trade-offs between aerodynamics and these other disciplines [125].

7.8 Conclusions

In this chapter, we have presented an extensive study of the CRM wing shape optimization benchmark defined by the Aerodynamic Design Optimization Discussion Group (ADODG). The drag coefficient is minimized for one flight condition with respect to 720 shape design variables, subject to lift, pitching moment, and geometric constraints, using grids with up to 28.8 M cells. The drag coefficient of the optimized design was reduced by 8.5% relative to the CRM baseline: from 199.7 counts to 182.8 counts, with a zero-grid spacing value of 181.9 counts. We implemented a multilevel optimization procedure that significantly reduced the total computational time.

The single-point optimized design exhibits a small thickness-to-chord ratio (3.3%) at the tip, which would incur a large structural weight penalty in a real wing. Thus, we performed an additional optimization that did not allow for thickness reduction. While the optimal drag increased by five counts relative to the nominal case, the

associated reduction in structural weight would likely offset this penalty in a real wing when considering the overall aircraft performance.

The multi-modality of the aerodynamic shape optimization problem was examined by starting optimizations from randomly generated initial geometries. All optimal wings had similar airfoil shapes, with an mean difference of 1.2 in. The variation of the merit function between the multiple local optima confirm that these points are indeed local minima, and indicate that the design space consists of a convex bowl with a small flat bottom that is multimodal. Based on our data, the minimum drag coefficient values were within 0.1 counts (0.05%), and the radius of this flat bottom seems to be about 1.6 in. Given these small differences, it does not seem worthwhile to put much effort into finding the global minimum for this problem.

We studied the effect of the design variables by varying the number of defining airfoil sections and the number of control points for each of those sections. Reducing the number of airfoil control points from 48 to 12 resulted in a 0.9-count drag increase. The total number of optimization iterations also reduced with the number of airfoil control points. The number of airfoil sections has a similar influence on the optimized drag. However, decreasing the number of airfoil sections while keeping the number of airfoil control points constant did not affect the overall computational cost in a significant way. We found that the optimization with 8 airfoil sections and 24 control points per section (192 design variables) provided the best trade-off: it increased the optimal drag by only 0.6 counts relative to the 720-variable case (15 airfoils with 48 points each), while requiring 40% fewer optimization iterations.

Finally, we performed a multipoint optimization of the CRM wing. This resulted in a more robust design than that of the single-point optimization, as evidenced by the enlarged contour of the 99% maximum ML/D . We also compared the contours of ML/D for the single-point baseline optimum and the multipoint optimum. Both the single-point and multipoint optimizations shifted the maximum ML/D toward the nominal flight condition.

This CRM wing aerodynamic shape optimization problem is a valuable benchmark for the wing design optimization community, and we hope that more researchers tackle this problem. The ADODG is also expected to expand this suite of benchmark problems in the near future.

This aerodynamic design optimization problem is limited to a fixed wing planform, but it is an excellent first step. To consider span and sweep, and to eliminate the explicit thickness constraints, it is necessary to consider the trade-offs between drag and structural weight, which has been done in an optimal way using aerostructural

optimization [59]. To take full advantage of the optimization, we should also include the flight conditions as design variables, but then we would have to solve an even more complex MDO problem that considers propulsion, mission analysis, and economics.

CHAPTER 8

Aerodynamic Design Optimization of an Adaptive Trailing Edge Wing

Adaptive morphing trailing edge wings have the potential to reduce the fuel burn of transport aircraft. In this chapter, we quantify the aerodynamic performance benefits of a morphing trailing using aerodynamic design optimization. The aerodynamic model solves the Reynolds-averaged Navier–Stokes equations with a Spalart–Allmaras turbulence model. A gradient-based optimization algorithm is used in conjunction with an adjoint method that computes the required derivatives. The baseline geometry is optimized using a multipoint formulation with 192 shape design variables. The drag coefficient is minimized subject to lift, pitching moment, geometric constraints, and a 2.5 g maneuver bending moment constraint. The trailing edge of the wing is optimized based on the multipoint optimized wing. The trailing edge morphing is parameterized using 90 design variables that are optimized independently for each flight condition. A total of 407 trailing edge optimizations are performed with different flight conditions to span the entire flight envelope. This chapter is organized as the follows: The trailing edge morphing optimization results are presented in Section 8.4. Then, we discuss the full morphing wing optimizations and comparison with morphing trailing edge in Section 8.5 and Section 8.6. We simulate a number of flight missions and quantify the fuel burn reduction with the adaptive morphing trailing edge in Section 8.7. This work is based on a previous paper presented by the author. [102]

8.1 Geometric Parametrization

The FFD geometry parametrization is similar to that in previous chapters shown in Figure 8.1. However, to simulate the trailing edge morphing, the last 5 chordwise

control points (shown in blue), which correspond to the last 45% of the chord, can move independently for each flight condition, thus providing sufficient degrees-of-freedom to alter the airfoil camber and spanwise twist distribution to simulate the morphing trailing edge similar to the FlexSys adaptive wing [40]. Because of the constant topology assumption of the FFD approach, and due to limitations in the mesh perturbation, the surface has to be continuous around the control surfaces, eliminating the elevator gap. Therefore, when the control surfaces deflect, there is a transition region between the control surface and the centerbody, similar to those studied in a continuous morphing wing [40].




Figure 8.1: The wing shape design variables are the z -displacement of 192 FFD control points (red and blue spheres). The trailing edge morphing design variables are only the blue control points.

8.2 Optimization Problem Formulation

All optimization cases perform lift-constrained drag minimization of the wing using the RANS equations. The baseline geometry and mesh is the same as that in Chapter 6 and 7. Since we need to perform hundreds of optimizations to optimize the trailing edge for each flight conditions, we use the L2 mesh to achieve a reasonable computational cost with sufficient accuracy. For simplicity, we only use the L2 mesh for the studies in this chapter.

8.2.1 Objective Function

The baseline multipoint aerodynamic shape optimization seeks to minimize averaged drag coefficients by varying the shape design variables subject to constraints on the lift, pitching moment, and maneuver bending moment. The drag coefficients

are given by the RANS solutions. Drag coefficients of 5 flight conditions are considered as shown in Table 8.1. The bending moment constraint is computed at a 2.5 g maneuver condition (15,000 ft and Mach 0.86). The formulation is different from the ADODG benchmark case. A similar multipoint optimization has been presented by the authors [101].

Flow Case	C_L	Mach number
1	0.50	0.85
2	0.55	0.85
3	0.45	0.85
4	0.50	0.84
5	0.50	0.86
6	2.5 g	0.86

Table 8.1: The multiple flight conditions represent a five-point stencil in Mach- C_L space and a 2.5 g maneuver case.

8.2.2 Design Variables

Before we study the trailing edge morphing, we performed a multipoint aerodynamic shape optimization of the wing to obtain an optimized aerodynamic performance of the wing itself. The first set of design variables consists of control points distributed on the FFD volume. A total of 192 shape variables are distributed on the lower and upper surfaces of the FFD volume, as shown in Fig. 6.2. The large number of shape variables provides more degrees of freedom for the optimizer to explore, and this allows us to fine-tune the sectional airfoil shapes and the thickness-to-chord ratios at each spanwise location. Because of the efficient adjoint implementation, the cost of computing the shape gradients is nearly independent of the number of shape variables [58]. The full morphing wing optimization uses the same set of shape design variables.

For the morphing trailing edge optimization, we use a subset of the shape control point near the TE as the design variables, as shown in blue in Figure 6.2. Only the shape on the last 45% of the chord is allowed to change. The shape of the forward wing remains constant.

8.2.3 Constraints

Since optimizers tend to exploit any weaknesses in numerical models and problem formulations, an optimization problem needs to be carefully constrained in order

to yield a physically feasible design. We performed a multipoint optimization with 6 flight conditions: 5 cruise conditions and a 2.5 g maneuver condition. Both lift and pitching moment are constrained at the nominal flight condition (Mach 0.85, $C_L = 0.5$). In addition, the wing root bending moment is constrained to be less or equal than the nominal value at the 2.5 g maneuver condition. We also implement several geometric constraints. First, we impose constant thickness constraints from the 1% chord at the LE to the 99% chord near the TE. A total of 750 thickness constraints are imposed in the 25 by 30 grid. The constraints have a lower bound of 100% of the baseline thickness and no upper bound. These constraints ensure sufficient height in the centerbody cabin and sufficient fuel volume. The LE thickness constraint allows for the installation of slats, and the TE thickness is limited due to manufacturing constraints. The total volume of the wing is also constrained to meet a fuel volume requirement. Complete optimization problem is described in Table 8.2.

	Function/variable	Description	Quantity
minimize	C_b	Drag coefficient	
with respect to	α	Angle of attack	1
	z	FFD control point z -coordinates	192
		Total design variables	193
subject to	$C_L = C_L^*$	Lift coefficient constraint	6
	$C_{M_y} \geq -0.17$	Moment coefficient constraint	1
	$C^{bend} \geq C^{bend*}$	Bending moment coefficient constraint	1
	$t \geq t_{base}$	Minimum thickness constraints	750
	$V \geq V_{base}$	Minimum volume constraint	1
		Total constraints	759

Table 8.2: Aerodynamic shape optimization problem

8.3 Baseline Multipoint Aerodynamic Shape Optimization of the Wing

Before we perform any morphing trailing edge optimization, we first optimize the wing itself using a multipoint formulation to achieve a reasonable performance that is robust with respect to different flight conditions. In this section, we present our aerodynamic design optimization results for the wing (described in Table 6.2) under the 5 flight conditions and a 2.5 g maneuver condition. We use the L2 grid (450k cells) for the optimization. Transport aircraft operate at multiple cruise conditions because of variability in the flight missions and air traffic control restrictions. Single-point optimization at the nominal cruise condition could overstate the benefit of the optimization, since the optimization improves the on-design performance to the detriment of the off-design performance. The single-point optimization benchmark problem developed by the ADODG resulted in an optimal wing with an unrealistically sharp leading edge in the outboard section of the wing [101]. This was caused by a combination of the low value for the thickness constraints (25% of the baseline) and the single-point formulation [115, 101]. Therefore, in this study, we use multipoint formulation and 100% thickness constraints, which we have found to result in more realistic wings [101].

We choose five equally weighted flight conditions with different combinations of lift coefficient and Mach number, as previously done by the authors [59, 101]. The flight conditions are the nominal cruise, $\pm 10\%$ of cruise C_L , and ± 0.01 of cruise Mach, as shown in Table 8.1. More sophisticated ways of choosing multipoint flight conditions and their associated weights can be used, such as the automated procedure developed by Liem *et al.* [125] that minimizes fleet-level fuel burn. The objective function is the average drag coefficient for the five flight conditions, and the moment constraint is enforced only for the nominal flight condition. The bending moment constraint is enforced at the 2.5 g maneuver condition at 15,000 ft and Mach 0.86.

A comparison of the initial wing and multipoint optimized design is shown in Figure 8.2. The baseline results are shown in red, and the multipoint results are shown in blue. The C_p for the multipoint optimized result corresponds to the nominal condition (Mach 0.85, $C_L = 0.5$). We compute the shock surface from the volume solution grid by constructing an isosurface of the normal Mach number [128]. The shock occurs where the normal Mach number is one, i.e.,

$$M_n = \frac{\vec{u}}{a} \cdot \frac{\nabla p}{|\nabla p|} = 1. \quad (8.1)$$




Figure 8.2: The multipoint optimized wing has 5.7% lower drag.

Unlike the shock-free design obtained with single-point optimization [101], the multipoint optimization settled on an optimal compromise between the flight conditions, resulting in a weak shock at all conditions. Similar trends were observed in the multipoint optimization of Vassberg *et al.* [116]. Our optimization procedure reduced the drag from 211.5 counts to 199.4 counts, i.e., an 5.7% reduction. At the optimum, the lift coefficient target is met, and the pitching moment is reduced to the lowest allowed value. The 2.5 g bending moment constraint is met. The lift distribution of the optimized wing is much closer to the elliptical distribution than that of the baseline, indicating an induced drag that is close to the theoretical minimum for planar wakes. This is achieved by fine-tuning the twist distribution and airfoil shapes. The baseline wing has a near-linear twist distribution. The optimized design has more twist at the root and tip, and less twist near mid-wing. This multipoint optimized wing provides a reasonable baseline geometry for the morphing trailing edge optimization.

8.4 Morphing Trailing Edge Optimization

We perform a series of RANS-based aerodynamic shape optimizations to examine the effects of trailing edge morphing. A gradient-based optimizer is used with sensitivities computed by the adjoint method. The full turbulence adjoint used in-

cludes the linearization of both the main flow solver and the SA turbulence model. The optimizations are converged to an optimality tolerance of $\mathcal{O}(10^{-5})$. We use the optimized geometry from the previous optimization in Section 8.3 as the baseline.

The shape aft of the 45% chord is free to change independently for each flight condition. The airfoil thickness is kept constant by the thickness constraints. A total of 80 design variables are used for each optimization. The angle-of-attack is also allowed to change during the optimization. To span the entire flight envelope, we performed 407 separate optimizations at various altitudes, Mach numbers, and weights. Each optimization required about 4 hours on 64 processors, corresponding to about 50 optimization iterations. No additional moment constraints are imposed in the optimization. Since the TE can be morphed at each flight condition, the 2.5 g maneuver bending moment constraints can be satisfied using the deflected TE. Therefore, we can check to see whether the 2.5 g bending moment constraint can be satisfied independently from the TE optimization at cruise conditions. Figures 8.3 to 8.6 show the trailing edge optimization results at several on- and off-design conditions. Results for additional flight conditions are shown in Figures C.1 to C.14.




Figure 8.3: Morphing trailing edge optimization at MTOW on-design condition.

At on-design conditions, the drag reductions range from 1 to 2%. The optimized TE shapes are extremely close to the initial shape. However, we see that the optimizer is able to further smooth out the flow by introducing a slight camber at the TE. The




Figure 8.4: Morphing trailing edge optimization at half-weight on-design condition.




Figure 8.5: Morphing trailing edge optimization at low-Mach low-altitude off-design condition.




Figure 8.6: Morphing trailing edge optimization at low-Mach high-altitude off-design condition.

TE deflection is less than 1 degree (measured from LE). The shock strength is reduced, illustrating that the transonic flow is sensitive to even slight changes in trailing edge shape. Similar trends are observed at several different weights.

At the off-design conditions, the difference between the optimized TE shape and the initial TE shape is more apparent. The maximum TE deflection at off-design conditions is about 3 degrees. The drag reduction due to morphing TE is more significant, reaching 5%. At some extreme off-design cases, the flow is separated without a morphing TE. By optimizing the TE camber and shape, the angle-of-attack is reduced and the flow is re-attached, significantly reducing the drag. The 2.5 g maneuver bending moment constraint is satisfied using the TE morphing. We conclude that the drag at all flight conditions can be reduced using morphing TE, and that the benefit of a morphing trailing edge is more significant at off-design conditions.

8.5 Full Wing Morphing Optimization

We also performed a shape optimization assuming a fully morphing wing. While the technology for achieving such morphing is currently not available, we are interested in finding out how much the performance would increase relative to the

morphing TE. A total of 192 design variables are used to optimize the entire wing at each flight condition. The rest of the optimization setup is the same as that in Section 8.4. Similarly, to span the entire flight envelope, we performed another 407 separate optimizations for different altitudes, Mach numbers, and weights. Due to the increased design space, the computational cost of the optimization is slightly higher: 6 hours on 64 processors instead of 4 hours in the TE morphing case. Figures 8.7 and 8.8 show the full morphing wing optimization results for an on-design and an off-design condition.

At on-design conditions, the full wing morphing designs are only marginally better than the optimized design with morphing TE. Specifically, the drag coefficient is decreased by about 1 count. The baseline wing is already optimized near the cruise conditions. Additional drag reduction is difficult to achieve even with a full morphing wing. The optimized wing shapes are very close to the initial shape. The pressure distributions are also quite similar to that of the morphing TE optimized designs. Therefore, we see that it is sufficient to only change the TE shape for drag reduction purpose at on-design conditions.

At the off-design conditions, additional improvements of up to 10 drag counts are achieved. The maximum TE deflection at off-design conditions is about 3 degrees. The drag reduction due to morphing TE is much more significant at more than 5%. In the flight condition shown in Figure 8.8, the flow on the initial wing is separated. The full morphing wing still maintains a shock-free solution and near-elliptical lift distribution even at high C_L . We observe that the benefit of morphing wing can be magnified at off-design conditions.

8.6 Comparison between Morphing Trailing Edge and Fully Morphing Wing

To further compare the benefits of the morphing TE and the morphing wing, we plotted the percentage drag reduction contours of each approach for the entire flight envelope for MTOW (347,500 kg), as shown in Figure 8.9 and 8.10. The drag reduction contours for other weights are shown in Figures C.15 to C.18. The weight and altitude range is based on the Boeing 777-200LR operation manual for Long Range Cruise (LRC).

The trends of both drag reduction contours are similar. The lowest drag reductions are near the on-design conditions where the wing has been previously optimized with a multipoint formulation. Those drag reductions are due to the additional degrees-of-




Figure 8.7: Full wing morphing optimization at MTOW on-design condition.

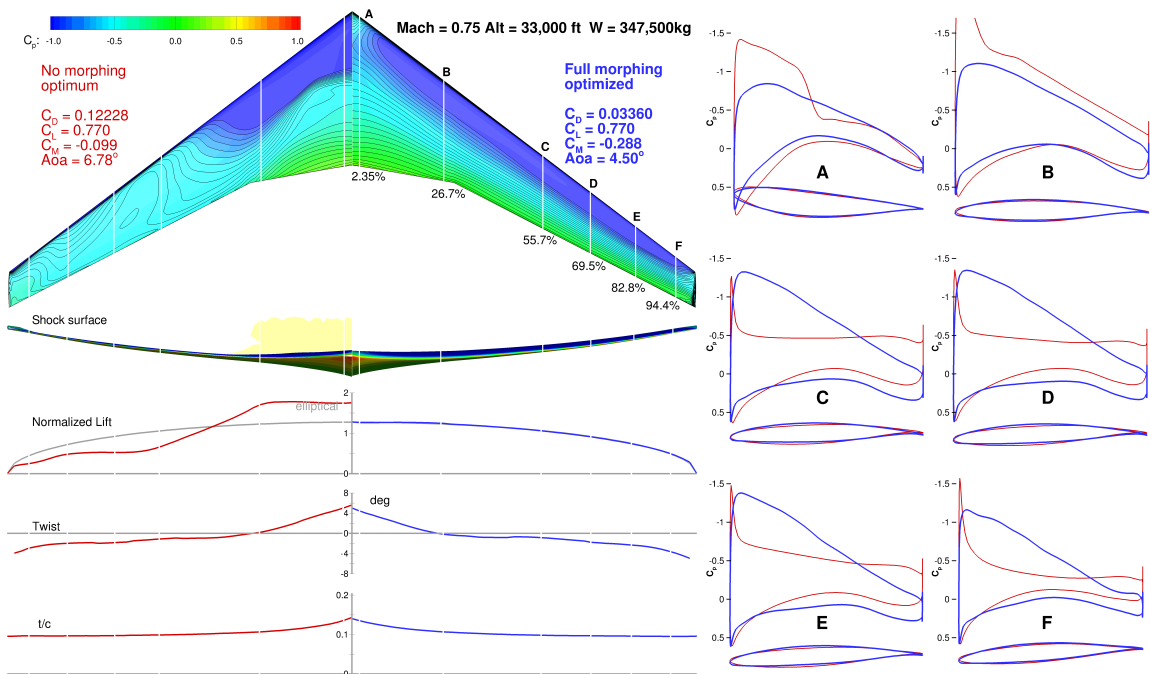


Figure 8.8: Full wing morphing optimization at low-Mach high-altitude off-design condition.




Figure 8.9: Drag reduction with morphing TE.




Figure 8.10: Drag reduction with full morphing wing.




Figure 8.11: ML/D contour of multipoint optimized baseline wing.




Figure 8.12: ML/D contour of the morphing TE wing.




Figure 8.13: ML/D contour of the full morphing wing.

freedom that allow the TE shape to change separately at each flight condition, and to making the 2.5 g maneuver condition constraint independent through load alleviation with the morphing TE. At the lower Mach number range, the drag reduction increases with the altitude and Mach number. The highest drag reduction occurred at the flight condition with high altitude and low Mach, where the lift coefficient is the highest. For high Mach numbers above 0.85, the trend reverses due to the drag divergence.

We also plot ML/D contours of the multipoint baseline, morphing TE, and fully morphing wing designs with respect to altitude and Mach number in Figure 8.11 to 8.13. ML/D contours for other weights are shown in Figures C.19 to C.24. ML/D provides a metric for quantifying aircraft range based on the Breguet range equation with constant thrust-specific fuel consumption. While the thrust-specific fuel consumption is actually not constant, assuming it to be constant is acceptable when comparing performance in a limited Mach number range [126]. We add 100 drag counts to the computed drag to account for the drag due to the fuselage, tail, and nacelles, and we get more realistic ML/D values. In aircraft design, the 99% value of the maximum ML/D contour, shown in black, is often used to examine the robustness of the design [37]. The point with the highest Mach number on that contour line corresponds to the long range cruise (LRC) point, which is the point at which

the aircraft can fly at a higher speed by incurring a 1% increase in fuel burn [127].

The multipoint baseline maximum ML/D occurs at the nominal flight condition (Mach 0.85, 31,000 ft altitude). Both morphing TE and morphing wing increase the maximum ML/D . The maximum ML/D points for morphing TE and morphing wing are at a higher altitude and higher Mach number. Since the TE shape can be adapted for each flight condition, the drag divergence is pushed to a higher Mach number. The 99% value of the maximum ML/D contour of the morphing designs are also significantly enlarged, indicating a more robust design. We see that morphing TE enables aircraft to fly higher and faster without a fuel burn penalty. To more accurately capture the tradeoffs, a multidisciplinary study including low speed aerodynamics, propulsion, and structure, would be required.

8.7 Simulating Flights with Morphing Trailing Edge

Since we have morphing trailing edge optimizations spanning the entire flight envelope, we can create a surrogate model of optimal trailing edge shapes for different flight conditions. This database allows us to compute the fuel burn for a series of missions without performing any additional optimizations. Since we have a relatively fine discretization of the flight region, we use a linear interpolation to evaluate the performance and optimal shape between the optimized points. A thrust specific fuel consumption (TSFC) of 0.53 lb/(lbf · h) is assumed. We also add 100 drag counts to the computed drag to account for the drag due to the fuselage, tail, and nacelles. The fuel burn is then integrated backwards for a given flight profile. Figure 8.14 shows a typical flight profile for a long range flight (currently the longest non-stop commercial flight from Dallas Fort Worth to Sydney, Australia).

Since the flight is operated in the on-design condition with step climb, the trailing edge deflection is within 1 degree. The wing tip exhibits the highest amount of deflection with -1 degree at the initial cruise to 1 degree near the end of the cruise. We see a 0.7% fuel burn reduction using morphing TE on this flight. As pointed out in Section 8.4, the morphing TE has higher drag reduction at off-design conditions. Table 8.3 shows the drag reduction on a number of hypothetical flight trajectories.

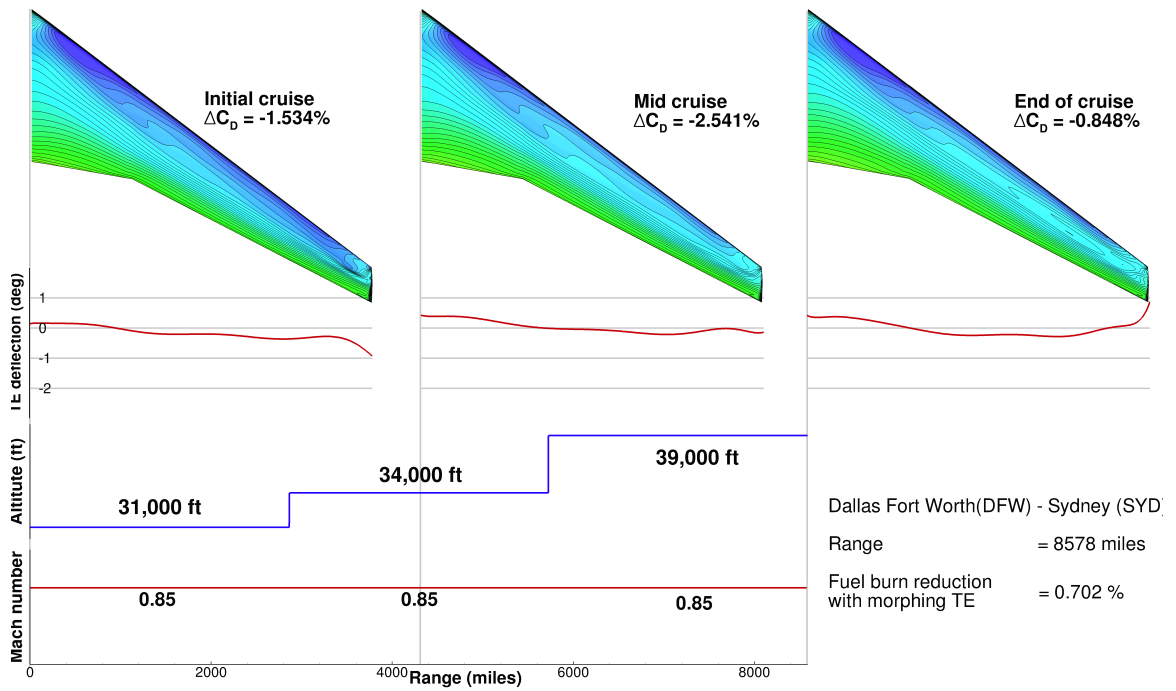


Figure 8.14: Fuel burn is reduced by 0.7% using morphing TE for DFW-SYD flight.

Departure	Arrival	Distance (miles)	Mach number	Cruise Altitude (feet)	Fuel burn reduction
DFW	SYD	8578	0.85, 0.85, 0.85	31,000, 34,000, 39,000	-0.702%
DTW	PVG	7120	0.85, 0.85, 0.85	33,000, 36,000, 41,000	-1.023%
LAX	NRT	5440	0.80, 0.85, 0.80	33,000, 33,000, 33,000	-1.049%
JFK	SFO	2580	0.84, 0.84, 0.84	32,000, 34,000, 36,000	-1.207%
ATL	ORD	606	0.78, 0.80, 0.78	31,000, 31,000, 31,000	-1.074%
-	-	2000	0.75, 0.75, 0.75	29,000, 31,000, 34,000	-1.68%

Table 8.3: The fuel burn reduction is about 1% using morphing TE for various flight trajectories.

We see that the morphing TE provides about 1% fuel burn reduction at cruise condition for the simulated flights in Table 8.3. All of the simulated flights have TE deflection within 2 degrees. Additional benefits could be realized during the climb and descent, which is neglected in this analysis. To evaluate the climb and descent, additional optimizations at lower speeds and lower altitudes would be needed to span the flight envelope for climb and descent.

8.8 Conclusions

In this chapter, we presented the aerodynamic shape optimization of a Boeing 777-size wing with an adaptive morphing trailing edge. A multipoint optimization, including a 2.5g maneuver condition, was presented to provide a baseline for the TE optimization. A total of 407 trailing edge optimizations with different Mach number, altitude, and weight, were performed to span the entire cruise flight envelope. A drag reduction in the order of 1% is achieved for on-design conditions, and reductions up to 5% were achieved for off-design conditions.

We further evaluated the performance of a morphing trailing edge by comparing its benefits with those from a full morphing wing. This is done by plotting the drag reduction contour and the ML/D contour. The full morphing wing yielded only marginally lower drag and a similar ML/D contour. Therefore, morphing only the TE can achieve an aerodynamic performance similar to that of a fully morphing wing without the drastic increase in wing morphing mechanism and weight.

Finally, we created a surrogate model of optimal trailing edge shapes to compute cruise fuel burn for different flight missions. We observed about 1% fuel burn reduction using the morphing trailing edge. More significant fuel burn reduction could be achieved in climb and descent segments.

From an aerodynamic perspective, an adaptive morphing trailing edge can easily offer additional drag reduction without a complete redesign of the wing. Since this technology has been demonstrated by FlexSys, and could be installed on conventional control surfaces, we could consider retrofitting existing aircraft. To thoroughly evaluate the benefit, a multidisciplinary study is required to examine the trade-offs between aerodynamics, structures, and controls.

CHAPTER 9

Aerodynamic Design Optimization of a Blended-Wing-Body Aircraft

The blended-wing body is an aircraft configuration that has the potential to be more efficient than conventional large transport aircraft configurations with the same capability. However, the design of the blended-wing is challenging due to the tight coupling between aerodynamic performance, trim, and stability. Other design challenges include the nature and number of the design variables involved, and the transonic flow conditions. With the aerodynamic shape optimization framework developed in this thesis, we can address these issues by performing a series of aerodynamic shape optimization studies using Reynolds-averaged Navier–Stokes computational fluid dynamics with a Spalart–Allmaras turbulence model. In this final chapter, we explore the potential of applying numerical optimization to an entire unconventional aircraft configuration.

The objective of this chapter is to develop a methodology for the aerodynamic design of BWB configurations that performs optimal trade-offs between the performance and constraints mentioned above, and to examine the impact of each constraint on optimal designs. We investigate the design trade-offs by performing a series of aerodynamic shape and planform optimization studies that examine the impact of the design variables and constraints. We explore the effect of the trim constraint, required static margin, and CG location on the BWB optimal shape. We also investigate the impact of multi-point design optimization. This work extends our preliminary studies to multi-point RANS-based aerodynamic shape and planform optimization [14]. The work in this chapter is based on previous papers presented by the author. [13, 14, 100]

9.1 Geometric Parametrization

Figure 9.1 shows the FFD volume and geometric control points for the BWB aerodynamic shape optimization. To trim the BWB configuration, we use control surfaces on the rear centerbody, which are analogous to elevators on a conventional configuration. A nested FFD volume is used to implement the movement of these control surfaces, as shown in Fig. 6.2. The result is a sub-FFD that is embedded in the main FFD. Any changes in the main FFD are propagated to the sub-FFD. The sub-FFD is set to rotate about the hinge line of the control surface. When the sub-FFD rotates, the embedded geometry changes the local shape accordingly. Because of the constant topology assumption of the FFD approach, and the limitation of the mesh perturbation, the surface has to be continuous around the control surfaces, eliminating the elevator gap. Therefore, when the control surfaces deflect, there is a transition region between the control surface and the centerbody, similar to those studied in a continuous morphing wing [40]. Figure 9.2 shows the sub-FFD volume and the geometry, with a trim control surface deflection of 25 degrees.




Figure 9.1: FFD volume (black) and control surface sub-FFD volume (red) with their respective control points

9.2 Problem Formulation

The BWB configurations can have more significantly improved aerodynamic performance than conventional configurations do. To fully realize this potential, however,




Figure 9.2: Sub-FFD volume and control points for a trim control surface deflection of 25 degrees

the external shape of the BWB has to be carefully designed. The primary focus of this study is drag minimization subject to a lift constraint. Additionally, we consider the following constraints: trim, static margin, and bending moment. In this section, we discuss the problem setup and the optimization formulation for the aerodynamic shape optimization of the BWB.

9.3 Initial Geometry

The initial geometry is shown in Fig. 9.3. The BWB geometry has a similar planform shape to the first-generation Boeing BWB design with 800 passengers [44]. This geometry has a span of 280 ft and a total length of 144 ft; it is divided into a centerbody section and an outer wing section. Based on this planform, the mean aerodynamic chord (MAC) is 86 ft. The initial CG is at 40% MAC of the planform. The placement of the CG is studied in Section 9.7.3.

The geometry is generated with a prescribed thickness-to-chord ratio (t/c), 18% at the center plane and 10% at the tip, as well as prescribed leading edge (LE) and trailing edge (TE) locations. We use the NASA SC(2)-0518 airfoil at the center plane and the NASA SC(2)-0410 airfoil at the tip, and we quadratically interpolate the airfoil sections in between. Table 9.1 summarizes the geometric parameters of the baseline BWB. The reference area is the actual area of the whole planform.

9.4 Grid Convergence Study

We generate the mesh for the BWB using an in-house hyperbolic mesh generator. The mesh is matched out from the surface mesh with an O-grid topology. The




Figure 9.3: Geometry of the BWB with the CG location shown in red

Geometric Parameter	Value
Span	280 ft
Length	144 ft
Reference area	15,860 ft ²
Mean aerodynamic chord	86 ft

Table 9.1: Geometric parameters for the BWB

nominal cruise flow condition is Mach 0.85 at 35,000 ft, and the Reynolds number is 100 million based on MAC. The spacing on the first layer uses a y^+ of 0.5 to adequately resolve the boundary layer. The grid is matched out to a far field that is located at a distance of 25 times the span, with an average growth ratio of 1.2. The grid used for the optimization has 2.92 million cells. It is generated from a surface mesh with 120 spanwise cells and 120 chordwise cells on each surface. There are also additional cells for the finite TE thickness and the rounded wingtip, resulting in a total of 30,464 surface cells. The resulting O-grid has 96 cells in the k direction.

We perform a grid convergence study to determine the resolution accuracy of this grid. All the grids are generated using the hyperbolic mesh generator with

a coarse or refined spacing. Figure 9.4 shows the mesh convergence plot, showing that the result for the mesh with 2.92 million cells is within 3 drag counts of that for the mesh with 187 million cells. We choose the former grid because it allows a reasonable optimization run time while providing sufficient accuracy. The RANS flow solution can be obtained within 100 minutes from a cold start with 6 orders of residual reduction on 180 processors. Figure 9.5 shows the BWB mesh on the surface and the symmetry plane.




Figure 9.4: Mesh convergence plot of the initial BWB mesh at nominal cruise condition

9.5 Optimization Problem Formulation

9.5.1 Objective Function

For the optimization studies, we minimize the drag coefficient at the nominal cruise condition, subject to a lift coefficient constraint. The drag coefficient is given by the RANS solutions. The cruise lift coefficient is constrained to $C_L = 0.206$. The chosen C_L is similar to that of the first-generation Boeing BWB [44], assuming a cruise altitude of 35,000 feet and a cruise Mach of 0.85. Since both the lift and drag coefficients use the whole planform area as the reference area, this results in a lower wing loading and lift coefficient.




Figure 9.5: BWB mesh showing surface and center plane cells

9.5.2 Design Variables

The first set of design variables consists of control points distributed on the FFD volume. A total of 240 shape variables are distributed on the lower and upper surfaces of the FFD volume, as shown in Fig. 6.2. The large number of shape variables provides more degrees of freedom for the optimizer to explore, and this allows us to fine-tune the sectional airfoil shapes and the thickness-to-chord ratios at each spanwise location. Because of the efficient adjoint implementation, the cost of computing the shape gradients is nearly independent of the number of shape variables [58].

The next set of design variables is the spanwise twist distribution. We use ten sectional twist design variables. The center of the twist rotation is fixed at the reference axis, which is located at the quarter chord of each section. The twist variables provide a way for the optimizer to minimize induced drag by controlling the spanwise lift distribution and a way to satisfy the center plane bending moment constraint.

We also consider planform variables, which can contribute to the reduction of wave drag. The sweep angle, chord length, and width of the centerbody are kept constant; only the planform variables of the outer wing are used as design variables. The outer wing is defined as the outer 60% of the total span, where the wing-centerbody

blending region ends. The outer wing is divided into seven sections. Each section has an independent set of planform variables, which are the sweep angle, chord length, and span of the section. Table 9.2 and Fig. 9.6 list the design variables. By providing complete freedom of the outer wing, we allow the optimizer to explore the optimal planform shape.

At the conceptual and preliminary design stages, the CG location should be optimized subject to trim and longitudinal stability constraints to minimize the trim drag. Thus, we use the CG location as a design variable that is allowed to move between 30% MAC and 50% MAC. In our case this variable represents the CG of the centerbody and the associated systems and payload. The CG of the wings is considered separately and is a function of the wing planform shape.

We add some auxiliary design variables to facilitate the formulation of the optimization problem. The angle-of-attack variable ensures that the lift coefficient constraint can be satisfied. We use an individual design feasible (IDF) approach [129] to update MAC. This requires the addition of a target variable and a compatibility constraint. With the IDF approach, the geometry manipulation and computation of MAC can be decoupled from the aerodynamic solver. Therefore, the sensitivity of MAC is also decoupled from the aerodynamic solver, which significantly simplifies the optimization problem formulation.

Design Variable	Count
shape	240
twist	10
sweep	7
chord	7
span	7
angle-of-attack	1
MAC ^t	1
Total	273

Table 9.2: Design variables for the BWB aerodynamic shape optimization

9.5.3 Constraints

Since optimizers tend to explore any weaknesses in numerical models and problem formulations, an optimization problem needs to be carefully constrained in order to yield a physically feasible design. We implement several geometric constraints. First,




Figure 9.6: Shape and planform design variables

we impose thickness constraints from the 5% chord at the LE to the 95% chord near the TE. A total of 400 thickness constraints are imposed in the 20 by 20 grid. The constraints have a lower bound of 70% of the baseline thickness and no upper bound. These constraints ensure sufficient height in the centerbody cabin and sufficient fuel volume. The LE thickness constraint allows for the installation of slats, and the TE thickness is limited due to manufacturing constraints.

The total volume of the centerbody and the wing is also constrained to meet the volume requirements for the cabin, cargo, and systems, as well as fuel. The LE and TE shape variables are constrained such that each pair of shape variables on the LE and TE can move only in opposite directions with equal magnitudes, so that twist cannot be generated with the shape design variables. Instead, twist is implemented as a separate set of variables.

Because of the absence of a structural model, we use the bending moment at the center plane as a surrogate for the structural weight trade-off and to prevent unrealistic spanwise lift distributions and wing spans. This bending moment is constrained to be less than or equal to the baseline bending moment. The bending constraint is necessary to capture the trade-offs between aerodynamic performance and structural weight. However, it is possible to perform these trade-offs with more accuracy by using high-fidelity aerostructural optimization, as done by Kenway and Martins [59].

In addition, the BWB has to be trimmed at each flight condition. Ideally, the

aircraft is trimmed at the nominal cruise condition without requiring control surface deflection. Therefore, we freeze the sub-FFD, which rotates the trim control surface during the on-design optimization with the pitching moment constraint. The sub-FFD is then used in the analysis of off-design conditions. There are several ways to trim a flying wing: by unloading wingtip on a swept wing, by adding reflex to the airfoils at the TE, or a combination of both of these [56]. Our optimization problem has all the required degrees of freedom to meet the trim constraint.

Longitudinal stability is also a particularly important design consideration for the BWB configuration. With the absence of a conventional empennage, it is not immediately obvious how to best achieve a positive static margin for a BWB aircraft. The goal is to maintain a positive static margin for all flight conditions. We constrained the static margin to be greater than 1%. The static margin, K_n , can be calculated as the ratio of the moment and lift derivatives [130, 131],

$$K_n = -\frac{C_{M_\alpha}}{C_{L_\alpha}}. \quad (9.1)$$

We calculate C_{M_α} and C_{L_α} using finite differences with an angle-of-attack step size of 0.1 deg. The static margin constraint incurs an additional computational cost. For each iteration, one additional flow solution and two additional adjoint solutions are required. Both the flow and adjoint solutions have to be converged more accurately than usual to obtain an accurate static margin gradient. This is particularly important for static margin gradients with respect to shape variables, because they have relatively small magnitudes compared to other gradients.

Table 9.3 summarizes the constraints for the optimization problems. All constraints are implemented as nonlinear constraints in the SNOPT optimizer.

Constraint	Count	Type
Thickness	400	<
LE, TE control points	40	<
Lift coefficient	1	=
Trim	1	=
Internal volume	1	<
Static margin	1	<
MAC compatibility	1	=
Total	445	

Table 9.3: Summary of the constraints used in the BWB aerodynamic shape optimization

9.6 Study 0: Baseline Optimization

To achieve a reasonable comparison for the optimization studies, we perform a baseline optimization by minimizing drag with respect to the spanwise twist distribution subject to a lift constraint. The airfoil profiles are the same as for the original geometry. The improved baseline has a drag 9 counts lower than that of the untwisted baseline. Studies in Section 9.7 use this improved baseline as the initial starting geometry for the optimization.

The improved baseline can still be improved upon, especially through changes in the sectional airfoil shape. Sensitivity of the drag and lift with respect to the airfoil shape can be visualized through a sensitivity contour plot, shown in Figs. 9.7 and 9.8. Here, we plot the derivatives of C_D and C_L with respect to shape variations in the y direction. The regions with the highest gradient of C_D are near the shock on the upper and lower surfaces. This indicates that shock reduction through local shape changes is the major driver in reducing C_D at the beginning of the optimization. As for C_L , a high positive gradient is observed near LE, indicating that moving in the positive y direction increases C_L . A high negative derivative is observed in the aft of the centerbody, indicating that moving the aft portion in the negative y direction increases C_L .

In addition, the regions with high derivative values on the lower and upper surfaces are offset longitudinally, which suggests that airfoil camber on the centerbody can further increase C_L . However, these sensitivity plots are only a linearization about the current design point, and they provide no information about the constraints. Nonetheless, these sensitivity plots indicate what drives the design at this design point.

9.7 Aerodynamic Design Optimization Studies of the BWB

We perform a series of RANS-based aerodynamic shape optimizations to examine the effects of various selections of design variables and constraints. The gradient-based optimizer (SNOPT) is used with sensitivities computed by the adjoint method. The full turbulence adjoint used includes the linearization of both the main flow solver and the SA turbulence model. Optimizations are converged to an optimality tolerance of $\mathcal{O}(10^{-5})$. By combining different sets of design variables and constraints, we explore the trade-offs and benefits of each. The initial design point for all the optimizations is the twist-optimized baseline described above. We then progress by




Figure 9.7: Study 0: dC_D/dy contour of the baseline BWB




Figure 9.8: Study 0: dC_L/dy contour of the baseline BWB

adding: airfoil shape variables (Study 1), a trim constraint (Study 2), a CG position variable and static margin constraint (Study 3), a bending moment constraint (Study 4), and planform design variables (Study 5). Finally, we consider multi-point opti-

mization (Study 6). This series of optimization studies allow us to examine how the optimization problem formulation impacts the practical design optimization of the BWB.

The optimizations are performed using the Advanced Research Computing cluster at the University of Michigan. Each computing node in this cluster has two six-core 2.67 GHz Intel Xeon X5650 processors per node. Each node has a total of 48 GB RAM. The cluster uses InfiniBand networking for interconnections.

9.7.1 Study 1: Shape and Twist Design Variables

In this first study, we add airfoil shape design variables to the twist variables already considered in the baseline optimization. A total of 240 shape design variables are used to optimize the airfoil shape. As shown in Fig. 9.6, 12 airfoil sections are equally distributed in the spanwise direction. Each section has 10 control points on the upper surface and 10 control points on the lower surface. The angle-of-attack is also allowed to change during the optimization. The CG is fixed at 40% MAC. Only lift and geometry constraints are imposed. Therefore, one flow solution and two adjoint solutions are needed at each iteration. The optimization converged in 10 hours using 240 processors; the convergence history of the optimization is shown in Fig. 9.9.

Without any additional constraints, we expect to see a lift distribution that is close to elliptical, along with weakened shocks. Figure 9.10 shows the pressure distribution, twist, sectional airfoil shape, shock surface, and lift distribution of the twist-optimized baseline and the optimized BWB for Study 1. A hypothetical elliptical lift distribution is shown in gray. We compute the shock surface from the volume solution grid by constructing an iso-surface of the normal Mach number [128]. The shock occurs where the normal Mach number is one, i.e.,

$$M_n = \frac{\vec{u}}{a} \cdot \frac{\nabla p}{|\nabla p|} = 1. \quad (9.2)$$

The dimensions in the figures are normalized by span, $\eta_s = z/2b$, and by chord, $\eta_c = x/c$.

As shown in the pressure distributions, the shape design variables make a significant contribution to the minimization of the drag. The baseline BWB exhibits a front of very closely spaced pressure contour lines spanning a significant portion of the wing, indicating a shock. The optimized BWB shows parallel pressure contour lines with roughly equal spacing, indicating a nearly shock-free solution at the nom-




Figure 9.9: Study 1: relative merit function (blue) and optimality (red) history of the optimization

inal cruise condition. This is confirmed by the shock surface plots: we can see that the baseline BWB has a shock on the upper surface, while the optimized design has eliminated most of the shock at the design condition. The shock elimination can also be seen on the airfoil C_p distributions. At $\eta_s = 0.4$ and $\eta_s = 0.9$, the sharp increase in local pressure due to the shock becomes a gradual change from the LE to the TE. The magnitude of C_p is also lowered near the LE.

The optimized lift distribution is much closer to the optimal elliptical lift distribution. This is achieved by altering the twist distribution. The highest twist is near the Yehudi break at $\eta_s = 0.6$, where the strong shock occurred on the baseline BWB. The drag coefficient is decreased by 39 counts. Twist angle at this section is increased to 5 degrees. The fact that the twist distribution has changed so much relative to a geometry that was already optimized for twist emphasizes the importance of simultaneously optimizing the twist and the airfoil shapes.

The angle-of-attack changed slightly from -0.4 deg to -1.0 deg. Since the CG is fixed at 40% MAC, the static margin is changed only by the shift in the aerodynamic center. The optimized design has reduced the static margin from 10.4% to 3.7%. A detailed study of the CG placement and static margin is presented in Section 9.7.3.

To study the effect of the twist variables, we performed a separate optimization

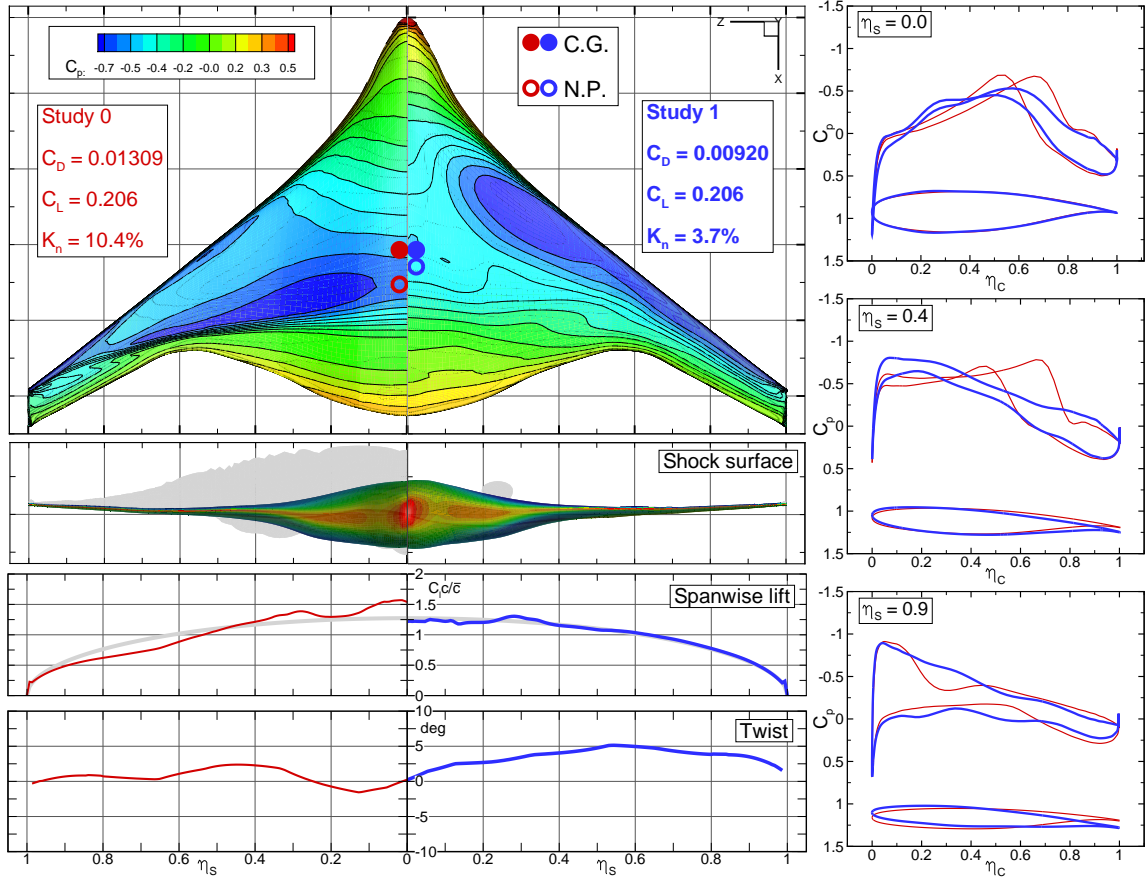


Figure 9.10: Study 1: optimized pressure distribution, sectional airfoil shape, shock surface, twist, CG, neutral point (NP), and lift distribution of Study 0 (red) and Study 1 (blue) BWB

that optimized only for the airfoil shape variables and did not include the twist design variables. The pressure distribution and airfoil profiles were similar to those for the case where both twist and airfoil shape were optimized. As shown in Table 9.4, the penalty for not including the twist variables is only 1.4 counts. Note, however, that we always start the optimization with the twist-optimized baseline geometry. Since the total variation of the optimized twist distribution is less than 5 degrees, the airfoil shape variables are able to get close to the overall optimum. Note that only aerodynamic performance is considered in this study. Other considerations, such as stall speed or wing structure, would pose additional constraints on the twist distribution.

Coefficient	C_D	C_L	AoA
Baseline	0.01309	0.206	-0.4
Twisted Optimum	0.00920	0.206	-1.0
Fixed Twist Optimum	0.00934	0.206	-0.2

Table 9.4: Study 1: comparison of twist design variables

9.7.2 Study 2: Trim Constraint

In this study, we investigate the effect of a trim constraint with a fixed CG. The formulation is the same as that of the previous study, with the addition of a trim constraint. Trim drag is more of a design driver in the BWB than in conventional tube-and-wing configurations, because elevator trim affects the flow around the BWB centerbody. Trim is also coupled to the longitudinal stability. Figure 9.11 shows the pressure distribution, twist, sectional airfoil shape, shock surface, and lift distribution of the optimized BWB for Studies 1 and 2.

The overall pressure contour is similar to that of Study 1. Compared to Study 1, the upper surface shock has increased at the optimum. However, it is still less severe than that of the baseline. The twist angles on both the centerbody and the outer wing are reduced. The wing has nearly zero twist for a large portion of the outer wing. The wingtip has a negative twist of 3 degrees (washout) in order to satisfy the trim constraint. Two design features helped satisfy the trim constraints of the optimized BWB. The first is a reflex near the TE throughout most of the span, resulting in a significant change to the chordwise pressure distribution. Angle-of-attack changed from -0.4 deg to 1.6 deg. Most lift is generated at the forward section of the wing, while the aft section has significantly less lift to trim the aircraft. Therefore, the net lift near the tip is reduced. The second feature is the unloaded wingtip. The optimized wingtip airfoil has washout and less lift than that of Study 1. Unloaded wingtip on a highly swept wing acts as a horizontal tail to trim the aircraft. Because of the trim constraint, the optimized drag coefficient is 5 counts higher than that of the previous study. This change is primarily due to a lower span efficiency and the reflex in the TE.

To investigate the off-design conditions, we perform a Mach sweep from 0.6 to 0.875. We use a sub-FFD to deflect the control surface near the rear centerbody to trim the aircraft at each condition, as shown in Fig. 6.2. The results are compared with the twist-optimized baseline design in Fig. 9.12. By comparing the trimmed baseline and optimized designs, we see that a trimmed drag pocket is achieved in the

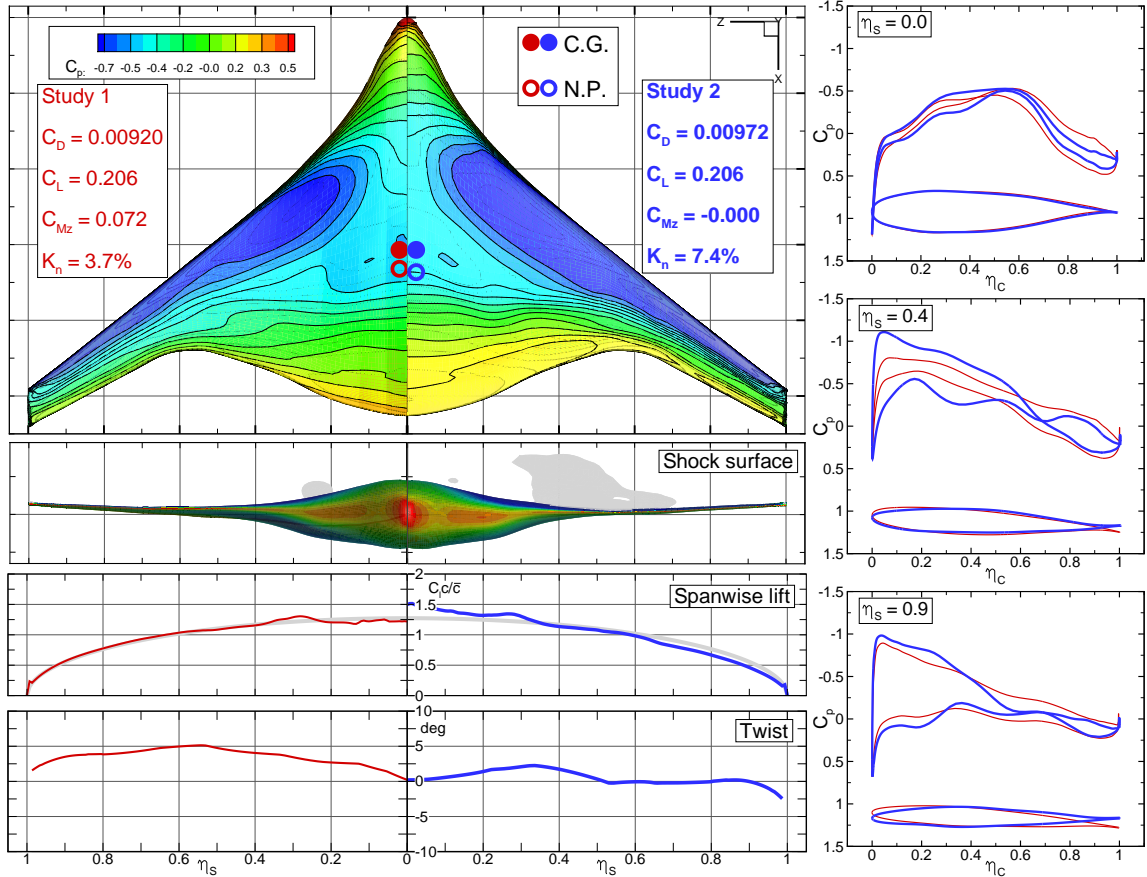


Figure 9.11: Study 2: optimized pressure distribution, sectional airfoil shape, shock surface, twist, CG, NP, and lift distribution of Study 1 (red) and Study 2 (blue) BWB

transonic region from Mach 0.80 to 0.86. The baseline design starts the drag rise near Mach 0.80, while the optimized design significantly delays the drag rise. The drag coefficient of the optimized design remains nearly constant up to Mach 0.86.

In addition, by comparing the trimmed and untrimmed results, we can quantify the trim drag at each condition. Figure 9.13 shows the trim drag of the baseline and optimized BWB. We see that the baseline design has lower trim drag at low Mach numbers. The optimized design, however, reverses this trend, and the trim drag reduces with increasing Mach number up to the design Mach number. Although this is a point design, the trim drag is relatively insensitive to the Mach number around the design point.




Figure 9.12: Mach sweep of trimmed and untrimmed designs for Studies 0 and 2




Figure 9.13: Trim drag of optimized BWB for Studies 0 and 2

9.7.3 Study 3: CG Design Variable and Static Margin Constraint

In the previous study, we examined the trim constraint with a fixed CG location. At the conceptual design stage, the CG can often be changed by moving systems, fuel, engines, and payload. By allowing the CG to change within a given range, we may discover additional benefits. To investigate the effect of CG location, we performed the same optimization as in Study 2 at various CG locations: 30%, 40%, and 50% MAC. The results are summarized in Table 9.5.

Coefficient	C_D	K_n	AoA
30% MAC	0.01032	19.6%	3.1
40% MAC	0.00972	7.4%	1.6
50% MAC	0.00941	-1.8%	1.4

Table 9.5: Study 3: comparison of optimized aerodynamic coefficients at various CG locations

Both the drag coefficient and the static margin are strongly affected by the CG location. Since a lower trim moment is required for an aft CG location, the trim constraint tends to move the CG back. We see that as the CG moves aft, the drag coefficient decreases, and the amount of reflex and washout is reduced. As the CG moves aft, the static margin decreases. For a flying wing, the location of the NP coincides with the aerodynamic center. For a fixed planform, the aerodynamic center varies only slightly with the airfoil shape variables. Therefore, the resulting static margin of the optimized design varies nearly linearly with the CG location.

Since the CG location is limited by both the trim and static margin, the problem formulation with CG design variables is a well-posed optimization problem. Simply adding CG design variables alone would result in the CG being as far aft as possible. Therefore, the CG design variable has to be added in conjunction with the static margin constraint. We perform another optimization with the CG position as a design variable and a static margin constraint. The CG is allowed to vary between 30% to 50% MAC. The static margin constraint has a lower bound of 1%.

This optimization problem is more computationally intensive than the previous cases for two reasons. First, each iteration requires two flow solutions and six adjoint solutions to obtain the static margin and its gradient. Second, the static margin gradient is a second-order derivative, since it is the gradient of the lift and moment coefficient gradients. Therefore, to achieve an accurate static margin gradient, both the flow and adjoint solutions must be converged to a higher tolerance, $\mathcal{O}(10^{-8})$. Fig-

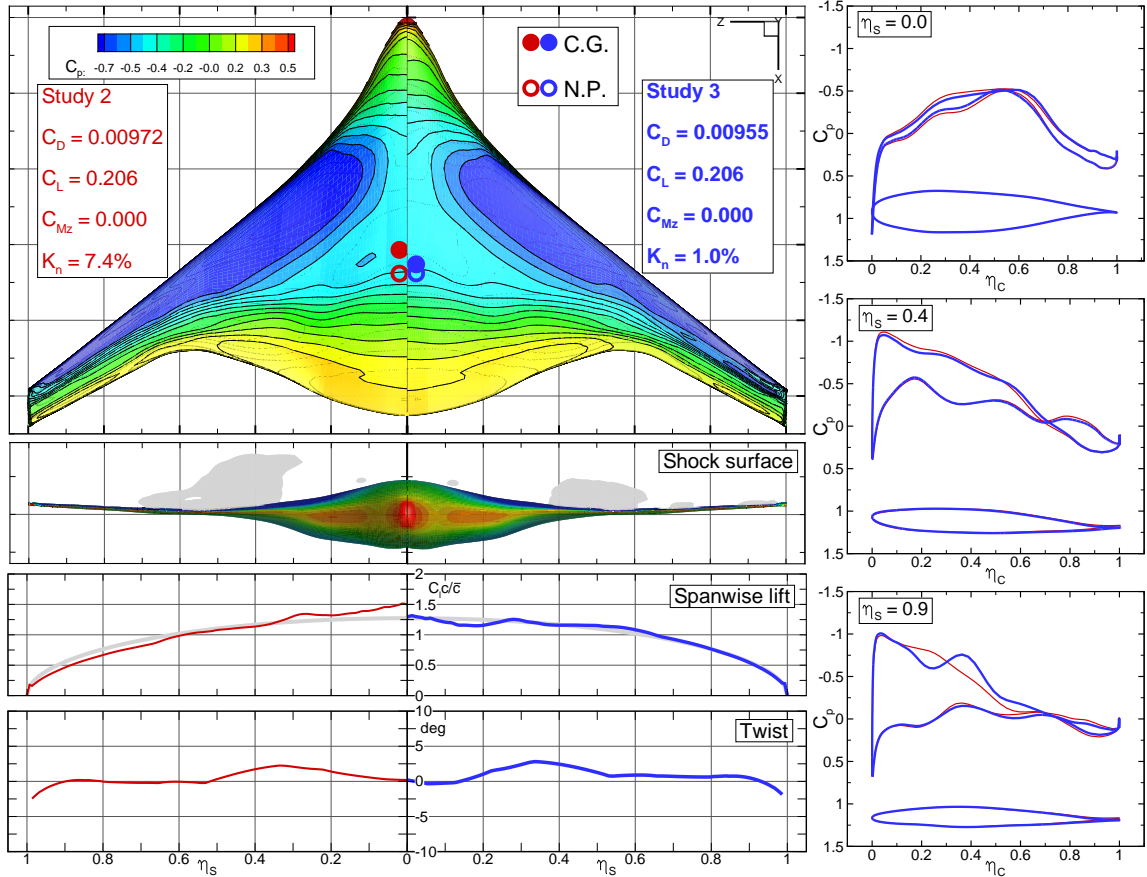


Figure 9.14: Study 3: optimized pressure distribution, sectional airfoil shape, shock surface, twist, CG, NP, and lift distribution of Study 2 (red) and Study 3 (blue) BWB

Figure 9.14 shows the pressure distribution, twist, sectional airfoil shape, shock surface, and lift distribution of the optimized BWB for Studies 2 and 3.

The overall pressure contours and airfoil profiles are similar to those for the optimal shape in the previous study. At the optimum, CG moves from 40% MAC to 47% MAC, driven by the trim constraint. Compared to Study 2, less airfoil reflex and wingtip unloading are needed to trim the BWB, resulting in an additional reduction of 1.7 in the drag count. The static margin is driven to the lower bound of 1%. In this study, CG is optimized based only on the aerodynamic performance and longitudinal stability. In reality, additional factors must be considered, such as the aircraft systems placement and the CG movement during operation, but these are beyond the scope of this study.

9.7.4 Study 4: Bending Moment Constraint

During the optimization, the aerodynamic load shifts. This may result in an increase in the structural stresses, which would impact the structural weight and thus the overall aircraft performance. A full aerostructural optimization, such as that presented by Kenway et al. [59], is beyond the scope of this work, but to limit the impact of the aerodynamic optimization on the structural weight, we add a center plane bending moment constraint [132]. This study is identical to Study 1 except for the addition of the bending moment constraint. The bending moment is taken about the center plane of BWB. We perform a series of optimizations with various bending moment constraints. The bending moment is constrained to be less than 100%, 80%, or 60% of the bending moment of the twist-optimized baseline. The results are summarized in Table 9.6.

Coefficient	C_D	C_{BM}	AoA
100% BM	0.00961	0.131	-1.8
80% BM	0.01103	0.105	0.9
60% BM	0.01399	0.078	3.5

Table 9.6: Study 4: comparison of optimized aerodynamic coefficients at various bending moment constraints

The addition of bending moment constraints drives the lift distribution away from elliptical. Figure 9.15 shows the lift distributions for each value of the bending moment constraint. A hypothetical elliptical span loading with the same lift is shown in gray. The optimization with the 100% bending moment constraint achieves a lift distribution that is the closest to elliptical. As the bending moment constraint decreases, more lift is shifted inboard to achieve the same lift with a reduced bending moment. A 20% reduction in the center plane bending moment results in a 14.2 increase in the drag count. A 40% bending moment reduction incurs a 43.8 increase in the drag count. At the reduced bending moment, the wingtip generates negative lift to alleviate the bending moment. Thus, we see that the impact of the bending moment constraint on aerodynamic performance is significant. For a careful trade-off between aerodynamics and structure, we would need to optimize both the aerodynamic shape and the structural sizing considering both the cruise performance and multiple load conditions [59].




Figure 9.15: Study 4: Spanwise lift distribution of optimized designs with various levels of bending moment constraint

9.7.5 Study 5: Planform Design Variables

In this study, we add planform variables to the previous study, which includes bending moment and trim constraints. The centerbody planform shape is kept constant. As shown in Fig. 9.6, the outer wing is divided into seven sections. Each section has its own twist, chord, sweep, and span design variables. The change in the planform shape, especially the span variables, would result in a heavier structure if no bending constraint were imposed. The center plane bending moment is constrained to be less than or equal to that of the twist-optimized baseline. Its CG is fixed at 40% MAC. MAC and the reference area are recomputed at each iteration to take the planform variations into account. The resulting optimized design is shown in Fig. 9.16. Outline of the baseline planform is shown in red.

The sweep angles of the outer wing of the optimized planform decrease by 4 degrees. The angle-of-attack of the optimized design is 0.6 deg. Even with the degrees of freedom provided by the multiple sweep, span, and chord of the outer wing sections, the optimization achieves a straight LE on the wing toward the end of the optimization. Any intermediate LE kinks during the optimization are smoothed out toward the end. The span of the optimized BWB increases by 3%. A further increase in the span to reduce the induced drag is constrained by the center plane bending moment and the additional viscous drag due to the increase in the surface area. Because of the presence of the trim constraint, the wing airfoil has reflex near

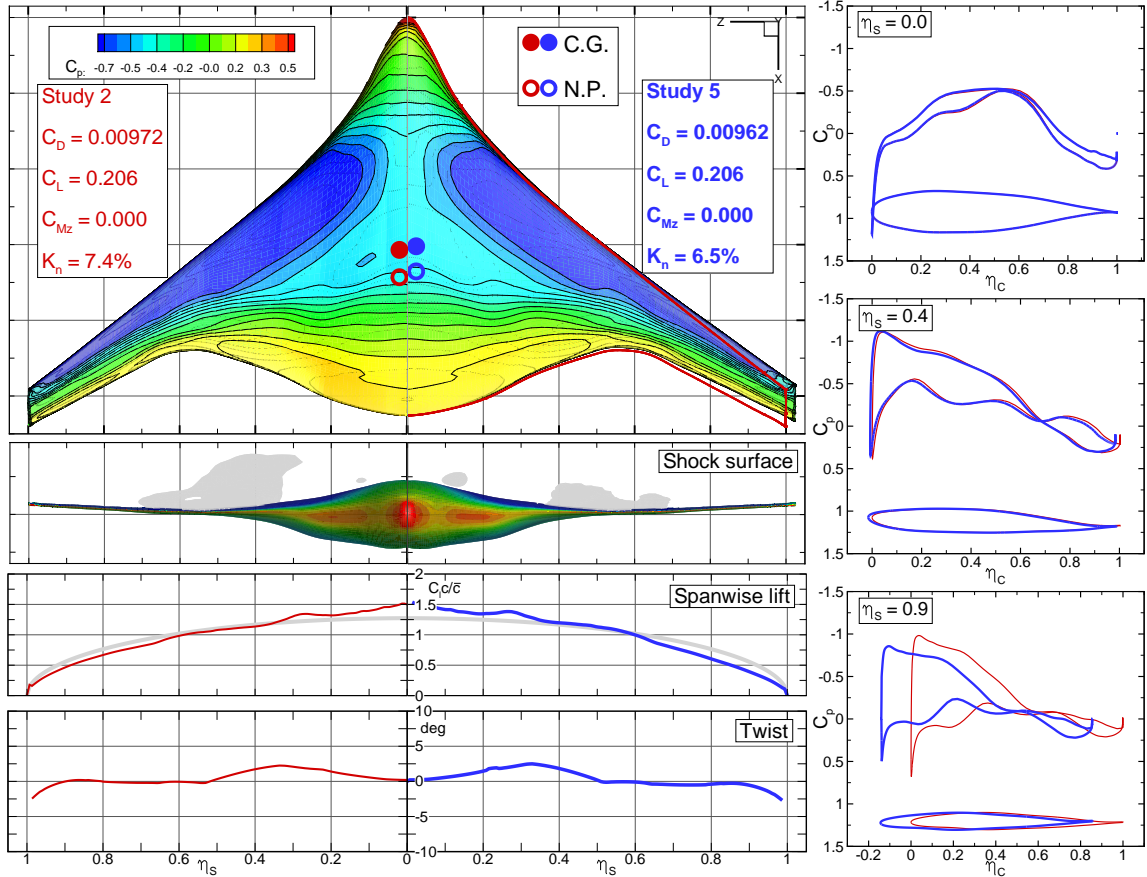


Figure 9.16: Study 5: optimized pressure distribution, sectional airfoil shape, shock surface, twist, CG, NP, and lift distribution of Study 2 (red) and Study 5 (blue) BWB

the TE, and the chordwise pressure distribution is similar to that of Study 2.

The planform study shows that the baseline is already relatively close to the optimal planform shape. Even with a marginal change in the planform, the additional degrees of freedom in the planform lead to a lower drag than that with only shape variables. The drag reduces by an additional drag count compared to Study 2, while satisfying the bending moment constraint.

9.7.6 Study 6: Multi-Point Optimization

Transport aircraft operate at multiple cruise conditions because of variability in both the missions and air traffic control restrictions. Single-point optimization at the nominal cruise condition could inflate the benefit of the optimization: it may improve the on-design performance while reducing the performance under off-design conditions. In this study, we investigate the impact of a multi-point optimization

formulation on the optimized BWB design. To isolate the problem from other effects of the constraints, we choose to extend Study 1. The only difference is that the objective is now the average of the drag coefficients at multiple flight conditions. The flight conditions are the nominal cruise, $\pm 10\%$ of cruise C_L , and ± 0.1 of cruise Mach, as shown in Table 9.7. More sophisticated ways of choosing multi-point flight conditions can be used, such as an automated selection of the points that minimize fleet-level fuel burn [133]. Figure 9.17 shows the multi-point optimized design at the nominal cruise condition. The multi-point optimized design is compared to the single-point optimum of Study 1.

Flight Condition	C_L	Mach
1	0.206	0.85
2	0.206	0.84
3	0.206	0.86
4	0.185	0.85
5	0.227	0.85

Table 9.7: Study 6: Flight conditions for the multi-point optimization

The overall pressure distribution of the multi-point design is similar to that of the single-point design. The twist and lift distributions are nearly identical. Most of the differences are in the chordwise C_p distributions in the outer wing section. Because of the multi-point formulation, the nominal cruise condition has less authority over the shape changes. The drag coefficient of the multi-point optimum is 2 counts higher than that of the single-point optimum, and the shock surface is also larger. Since all the flight conditions are equally weighted, the optimizer trades off the drag between the multiple flight conditions. The angle-of-attack of the optimized design at nominal flight condition is -0.6 deg.

To better understand the effects of multi-point optimization, we plotted the ML/D contours of the baseline, single-point, and multi-point designs with respect to C_L and cruise Mach in Figure 9.18 to 9.21. The line along which the aircraft is neutrally stable is shown in gray. ML/D provides a metric for quantifying aircraft range based on the Breguet range equation with constant thrust specific fuel consumption. While the thrust specific fuel consumption is actually not constant, assuming it to be constant is acceptable when comparing range performance in a limited Mach number range [126].

The baseline maximum ML/D is at a lower Mach number and a higher C_L compared to the cruise flight condition. The single-point optimization significantly in-

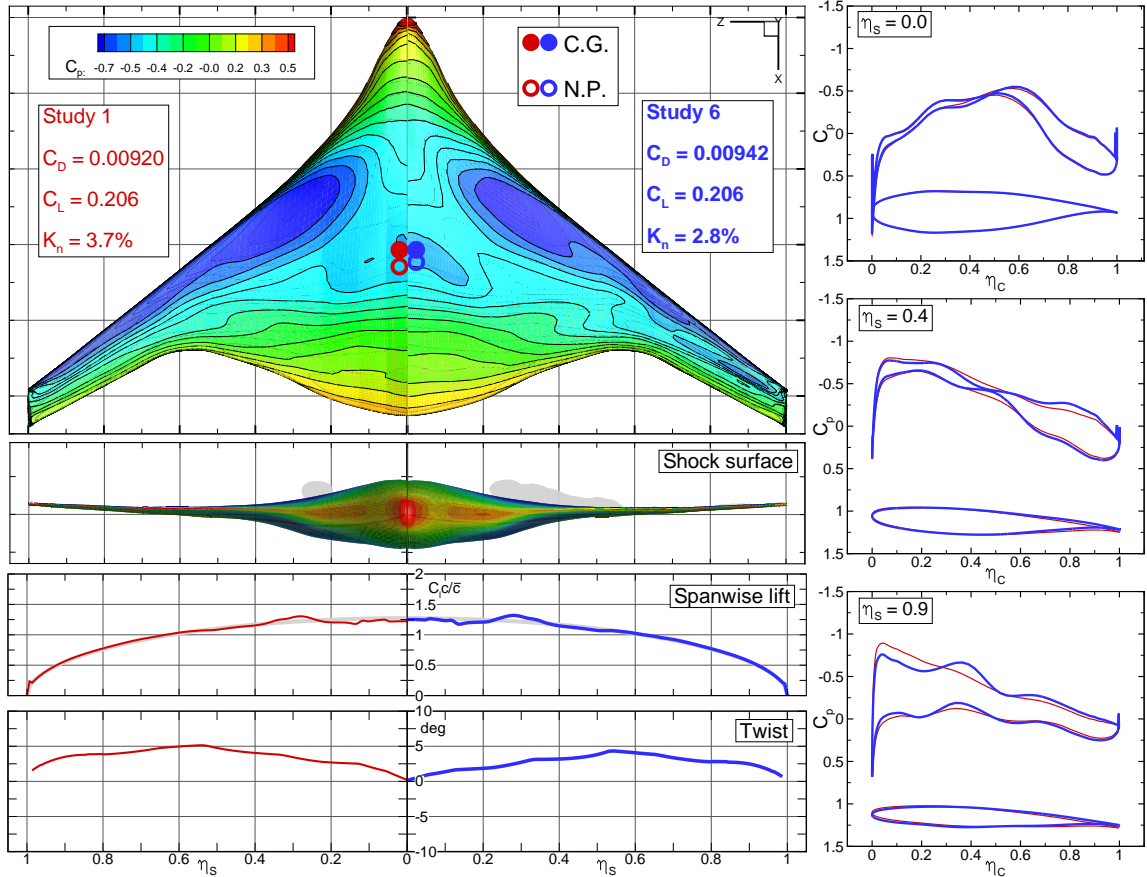


Figure 9.17: Study 6: optimized pressure distribution, sectional airfoil shape, shock surface, twist, CG, NP, and lift distribution of Study 1 single-point (red) and Study 6 multi-point (blue) optimized BWB

increases the maximum ML/D and the ML/D at the operation condition. In addition, the maximum ML/D occurs much closer to the nominal cruise condition. The shapes of the contours are also altered to move the maximum toward the cruise flight condition. For fixed $C_L = 0.206$, the maximum ML/D occurs near a cruise Mach of 0.85, which is equivalent to the drag bucket in a drag divergence plot. For the multi-point optimization, the flight conditions for optimization are spread in the Mach, C_L space, resulting in a more flattened ML/D near the maximum. ML/D is more uniform near the operation flight conditions. The 99% ML/D contour is also larger than that of the single-point optimum. By examining the ML/D , we see that the C_L of the maximum ML/D is still higher than the C_L in our optimization. An increase in the C_L may further improve the aerodynamic performance of the optimized BWB. The optimum C_L occurs between 0.25 and 0.27. Since the wing loading is constrained by the low speed performance, the only viable way to increase C_L for the BWB is to increase the cruise altitude. However, additional trade-offs, such as cabin pressure




Figure 9.18: Study 6: ML/D contours, 99% ML/D , and neutral stability line of twist-optimized baseline BWB




Figure 9.19: Study 6: ML/D contours, 99% ML/D , and neutral stability line of single-point optimized BWB




Figure 9.20: Study 6: ML/D contours, 99% ML/D , and neutral stability line of multi-point optimized BWB




Figure 9.21: Study 6: ML/D contours, 99% ML/D , and neutral stability line of all designs

and required thrust, must be taken into consideration.

9.8 Conclusions

In this chapter, we have presented a series of RANS-based aerodynamic shape optimization studies of a BWB configuration to understand the trade-offs between aerodynamic performance, and constraints on trim, stability, and bending moment. These studies also explored the effect of considering different sets of aerodynamic shape design variables (twist, airfoil shape and planform shape) in the design optimization. Table 9.8 summarizes the results of the optimization studies.

Study	C_D	Design Variables	CG	Constraints					
				K_n	Geo	Lift	BM	Trim	K_n
0	0.01309	twist	40%	10.4%	•	•			
1	0.00932	shape	40%		•	•			
	0.00920	shape, twist	40%	3.7%	•	•			
2	0.01032	shape, twist	30%	19.6%	•	•			•
	0.00972	shape, twist	40%	7.4%	•	•			•
	0.00941	shape, twist	50%	-1.8%	•	•			•
3	0.00955	shape, twist, CG	47%	1.0%	•	•			•
4	0.00961	shape, twist	40%		•	•	100%		
	0.01103	shape, twist	40%		•	•	80%		
	0.01399	shape, twist	40%		•	•	60%		
5	0.00962	shape, twist, planform	40%	6.5%	•	•	•		•
6	0.00942	shape, twist, multi-point	40%	2.8%	•	•			

Table 9.8: Summary of the results of BWB aerodynamic design optimization studies

The BWB configurations obtained in Studies 0 and 1 had the lowest drag coefficient, but they are impractical since they are not trimmed. The airfoil shape design variables proved essential to the reduction of the shock on the upper surface and the wave drag associated with it. Enforcement of the trim constraint (Study 2) caused a 5.6% increase in drag. By moving the CG aft from 40% MAC to 50% MAC, this drag penalty was reduced to 2.3%, but resulted in a negative static margin (-1.8%).

Study 3 provides the best compromise between performance and stability by enforcing a small static margin that can be tolerated in a commercial airplane (1%) and including the CG position as a design variable. This resulted in a trimmed configuration that exhibits a nearly elliptical lift distribution and the lowest drag among the

trimmed stable designs. This was achieved by a combination of washout and reflex airfoils determined by the optimizer to be the best.

We also investigated the optimized BWB at off-design flight conditions by analyzing it for a range of Mach numbers while enforcing trim. The optimized design exhibited significantly lower drag over the entire transonic regime when compared to the baseline. In the optimized design, we observed a low trim drag at high speeds and a high trim drag at low speeds, which is the opposite of the baseline BWB trend.

In Study 5 we further explored the design space by adding wing planform design variables to the optimization, while enforcing a center plane bending moment constraint. The addition of planform variables achieved an additional drag reduction. The optimized design increased the span by 3% and reduced the sweep angle by 4 degrees. This demonstrated the benefit of simultaneously optimizing the planform and shape, and highlighted the importance of aerostructural considerations. One of these considerations is the structural weight increase incurred by increases in span or sweep, which we addressed by enforcing the bending moment constraint. Since the right value for constraint requires a full aerostructural optimization that is beyond the scope of this thesis, we investigated the effect of varying the bending moment constraint on the optimal designs in Study 4. The results showed that when the bending moment constraint was reduced to 60% of the baseline value, the optimal design exhibited negative loading at the wingtips.

Finally, we studied the effect of multi-point optimization in Study 6. This resulted in a more robust design than that of the single-point optimization, as evidenced by the enlarged contour of the 99% maximum ML/D . We also compared the contours of ML/D for the twist-optimized baseline, single point optimum and multi-point optimum. These contours showed that the maximum ML/D occurs at a lower cruise Mach number and higher C_L than the design flight conditions, indicating that the configuration should either fly higher, or have a smaller planform area (although the engine performance would degrade with an altitude increase, and the planform area is probably constrained by field performance). Nevertheless, the aerodynamic shape optimization successfully moved the ML/D peak towards the design points, and flattened the peak in the multi-point case.

Given the results of these studies, we believe that RANS-based aerodynamic shape optimization has become a practical aircraft design tool that is especially useful for the design of BWB configurations. This type of optimization was enabled by the combination of a nonlinear constrained optimizer and an efficient computation of the gradients of the aerodynamic force coefficients with respect to hundreds of shape de-

sign variables. In the case of the BWB in particular, the optimal combination of wing twist, and airfoil reflex to obtain the lowest drag while satisfying trim, stability and structural constraints is not obvious, but numerical optimization can help designers find the best possible configuration.

CHAPTER 10

Final Remarks

This chapter presents the overall conclusions of this thesis, main contributions, and future research directions.

10.1 Conclusions

With recent improvements in the efficiency and fidelity of numerical simulations, aircraft design has become increasingly reliant on computational simulations and optimization. One of the most computationally intensive disciplines is the aircraft external aerodynamic design. Computational fluid dynamics based on Reynold-averaged Navier–Stokes equations is needed to accurately resolve the flow field in order to achieve a practical design. High-fidelity CFD poses difficulties to numerical optimization due to its high computational cost, especially when large number of shape design variables are used. Traditional optimization techniques, such as finite-difference derivatives, design of experiments, or evolutionary optimization methods, become less effective or even infeasible in high-fidelity aerodynamic shape optimization.

In Chapter 5, we demonstrated that gradient-based aerodynamic shape optimization is required in order to maintain reasonable computational cost. It is, however, often difficult to develop an efficient gradient computation to realize the benefit of gradient-based optimization. In Chapter 3, we presented an approach to compute the gradients of RANS equations with an SA turbulence model, using a combination of adjoint method and automatic differentiation algorithms. In addition, we developed a coloring acceleration technique to further improve the efficiency of gradient computations. We verified in Chapter 4 that the resulting gradients are accurate, robust, and efficient. We performed a RANS-based and Euler-based aerodynamic shape optimization of the ONERA M6 test case to demonstrate the need for RANS adjoint. We found that due to the lack of viscous effects, Euler-based optimization

tends to converge to non-physical optima, such as a design with rapid pressure recovery near the trailing edge of the wing. Even though RANS-based aerodynamic shape optimization is more computationally intensive, it achieves more practical designs and the optimization itself is also more robust. We also made improvements to the aerodynamic shape optimization framework by developing a novel multilevel optimization approach to reduce the computational cost by 84.5% in Chapter 6.

With this state-of-the-art RANS adjoint and aerodynamic shape optimization framework, we performed several high-fidelity aerodynamic design optimization studies in this thesis. In Chapter 7, we began by applying this aerodynamic shape optimization approach to improve the wing of a current-generation aircraft with a similar size to the Boeing 777-200, the NASA Common Research Model. The drag coefficient is minimized for one flight condition with respect to 720 shape design variables, subject to lift, pitching moment, and geometric constraints, using grids with up to 28.8M cells. The drag coefficient of the optimized design was reduced by 8.5% relative to the CRM baseline: from 199.7 counts to 182.8 counts, with a zero-grid spacing value of 181.9 counts. We also performed a multipoint optimization of the CRM wing. This resulted in a more robust design than that of the single-point optimization, as evidenced by the enlarged contour of the 99% maximum ML/D . We also compared the contours of ML/D for the single-point baseline optimum and the multipoint optimum. Both the single-point and multipoint optimizations shifted the maximum ML/D toward the nominal flight condition. In addition, the multi-modality of the aerodynamic shape optimization problem was examined by starting optimizations from randomly generated initial geometries. All optimal wings had similar airfoil shapes, with an mean difference of 1.2 in. The variation of the merit function between the multiple local optima confirm that these points are indeed local minima, and indicate that the design space consists of a convex bowl with a small flat bottom that is multimodal. Based on our data, the minimum drag coefficient values were within 0.1 counts (0.05%), and the radius of this flat bottom seems to be about 1.6 in. Given these small differences, it does not seem worthwhile to put much effort into finding the global minimum for this problem.

In Chapter 8, we extended the RANS-based aerodynamic shape optimization studies to a near-term aircraft retrofit modification, a wing with morphing trailing edge. A multipoint optimization, including a 2.5g maneuver condition, was presented to provide a baseline for the TE optimization. A total of 407 trailing edge optimizations with different Mach number, altitude, and weight values were performed, to span the entire cruise flight envelope. A drag reduction on the order of 1% is achieved

for on-design conditions, and reductions of up to 5% were achieved for off-design conditions. We further evaluated the performance of a morphing trailing edge by comparing its benefits with those from a full morphing wing. This is done by plotting the drag reduction contour and the ML/D contour. The full morphing wing yielded only marginally lower drag and a similar ML/D contour. Therefore, morphing only the TE can achieve an aerodynamic performance similar to that of a fully morphing wing, without the drastic increase in wing morphing mechanism and weight. Finally, we created a surrogate model of optimal trailing edge shapes to compute cruise fuel burn for different flight missions. We observed about 1% fuel burn reduction using the morphing trailing edge. More significant fuel burn reduction could be achieved in climb and descent segments.

We further the aerodynamic design optimization to an unconventional aircraft configuration, the blended-wing-body aircraft, in Chapter 9. The RANS adjoint and high-fidelity aerodynamic shape optimization framework allowed us to examine the trade-offs between drag coefficient, trim, and static margin of the BWB configuration. The airfoil shape design variables proved essential to the reduction of the shock on the upper surface and the wave drag associated with it. The enforcement of the trim constraint caused a 5.6% increase in drag. By moving the CG aft from 40% MAC to 50% MAC, this drag penalty was reduced to 2.3%, but resulted in a negative static margin (-1.8%). The best compromise between performance and stability was achieved by enforcing a small static margin that can be tolerated in a commercial airplane (1%) and including the CG position as a design variable. This resulted in a trimmed configuration that exhibits a nearly elliptical lift distribution and the lowest drag among the trimmed stable designs. This was achieved by a combination of washout and reflex airfoils determined by the optimizer to be the best. Given the results of these studies, we believe that RANS-based aerodynamic shape optimization developed in this thesis has become a practical aircraft design tool that is especially useful for the design of aircraft configurations. This type of optimization was enabled by the combination of a nonlinear constrained optimizer and an efficient computation of the gradients of the aerodynamic force coefficients with respect to hundreds of shape design variables. We believe that the numerical optimization can help designers find the best possible configuration.

10.2 Thesis Contributions

The optimization approaches and results presented in this thesis represent the current state-of-the-art in high-fidelity gradient-based aerodynamic shape optimization. The main contributions of this thesis are as follows:

1. Developed an accurate, efficient, and robust adjoint gradient algorithm based on Reynolds-averaged Navier–Stokes equations with a Spalart–Allmaras turbulence model. Forward mode automatic differentiation is used to reduce the development time for forming the partial derivatives in the adjoint formulation.
2. Developed coloring acceleration scheme for the RANS state and spatial Jacobians to allow efficient sparse matrix evaluations.
3. Implemented the first Euler and RANS-based aerodynamic optimization comparison to quantify the difference in the optimized designs.
4. Benchmarked several gradient-based and gradient-free optimizers on the aerodynamic shape optimization problems.
5. Developed a multilevel aerodynamic shape optimization acceleration technique to reduce the computational cost by 85%.
6. Performed RANS-based aerodynamic shape optimization of a transonic wing. The effects of shape design variables, thickness constraints, and single versus multi-point formulation were quantified.
7. Investigated the multi-modality of an aerodynamic shape optimization problem. Recommendation for future gradient-based aerodynamic shape optimization was provided.
8. Developed a series of optimal trailing edge shapes for a morphing trailing edge wing to reduce the fuel burn for an conventional aircraft. The aerodynamic benefit of morphing wing is evaluated and quantified.
9. Performed systematic RANS-based aerodynamic design optimization studies of the blended-wing-body aircraft. The trade-offs between aerodynamic performance, trim, and static margin were studied. Demonstrated the effectiveness of RANS-based aerodynamic shape optimization on an unconventional configuration.

10. Implemented planform design variables, such as sweep, span, and chord, on the blended-wing-body configuration. A planform optimization considering trim and stability of BWB was demonstrated.

10.3 Recommendations

Throughout this thesis, several research directions have been identified to further improve the RANS adjoint and advance the state-of-the art of high-fidelity aerodynamic shape optimization and aircraft design.

1. Although the forward mode RANS adjoint presented in this thesis is computationally efficient, the memory usage can be further reduced by computing $(\partial\mathcal{R}/\partial x_{pt})^T \Psi$ with reverse mode automatic differentiation and storing in a matrix-free fashion. A significant amount of memory storage can be reduced by only storing the matrix-vector product instead of the full Jacobian. The corresponding increase in computational time is expected to be moderate.
2. The linearization of the Spalart-Allmaras turbulence model is presented in the thesis. However, the automatic differentiation adjoint is not limited to any specific turbulence model. Linearizations of additional turbulence models can be implemented easily. A systematic study of effects of turbulence models on optimized designs can be beneficial to aircraft design engineers.
3. Another possible research direction is to extend the adjoint to the time domain. An implementation of a time-dependent adjoint could open the door to unsteady aerodynamic optimization. Related applications include the optimization of: helicopter blades, flapping wings, noise, etc.
4. In terms of the applications of the RANS adjoint, further improvements can be made to include more practical objectives and constraints to the optimization, such as low-speed constraints, flutter constraints, and structural constraints. In addition, the aerodynamic shape optimization can be extended to RANS-based aerostructural optimization to include aeroelasticity effects rather than optimizing a rigid wing shape.
5. Finally, the aerodynamic shape optimization framework can be used to explore additional aircraft configurations, such as the truss-braced wing configuration, double-bubble configuration, etc.

APPENDICES

APPENDIX A

Additional Data from Optimizers Comparison Study

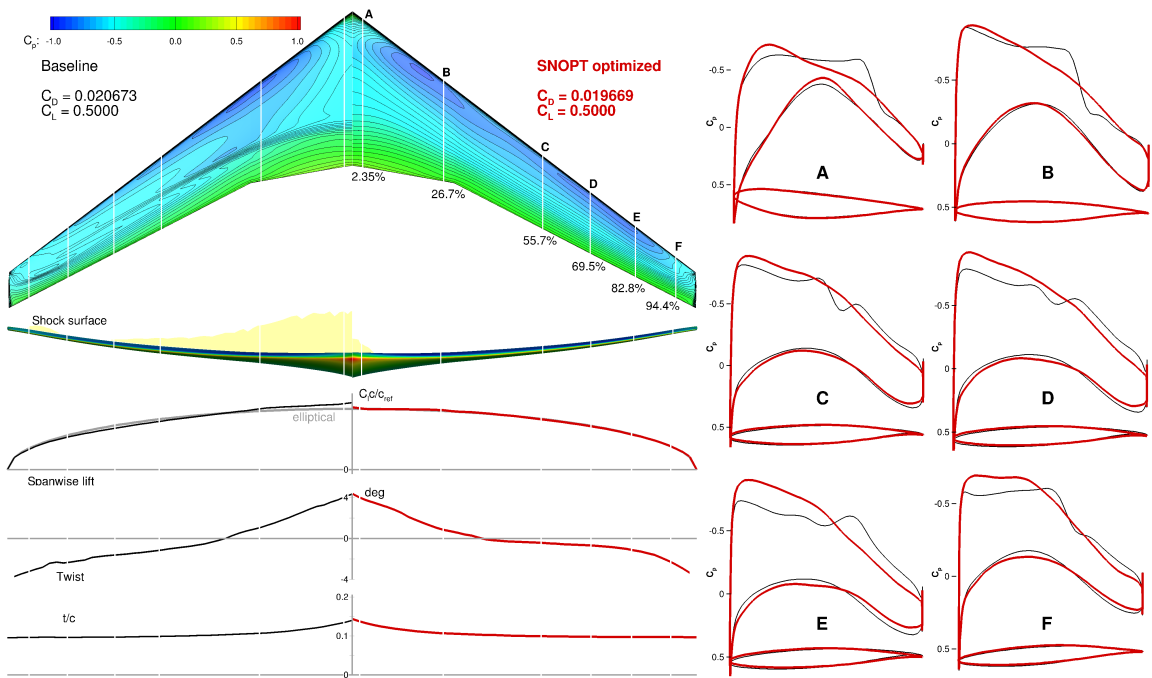


Figure A.1: Wing shape optimization using SNOPT




Figure A.2: Wing shape optimization using SLSQP

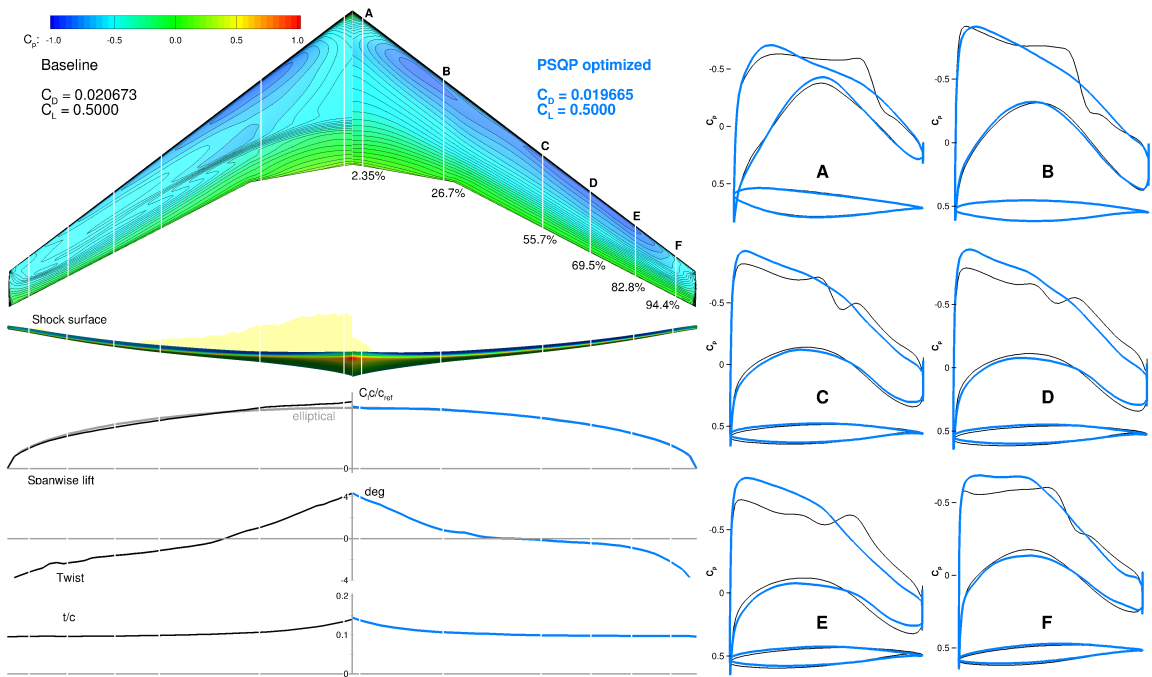


Figure A.3: Wing shape optimization using PSQP

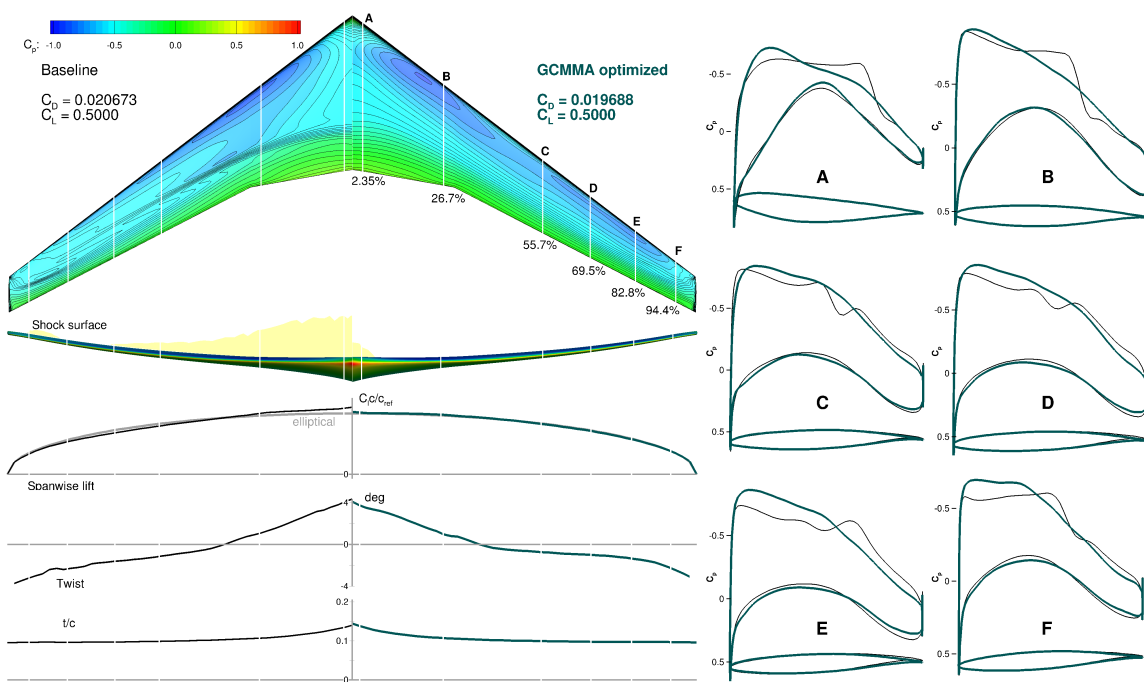


Figure A.4: Wing shape optimization using GCMMA

APPENDIX B

Comparison of ML/D , \sqrt{ML}/D , and $aM/c L/D$

In this appendix, we present the comparison of ML/D , \sqrt{ML}/D , and $aM/c L/D$ for the purpose comparing the performance and robustness of aircraft configurations.

B.1 Background

The Breguet range equation, shown in Equation B.1, describes the maximum range of a long range jet aircraft operating in the stratosphere.

$$R = \frac{aM}{c} \frac{L}{D} \ln \frac{W_1}{W_2} \quad (\text{B.1})$$

a is the speed of the sound. M is the Mach number. c is the thrust-specific fuel consumption rate (TSFC). W_1 is the weight of the aircraft at the beginning of the cruise, and W_2 is the weight of the aircraft at the end of the cruise. The speed of sound a in this equation is assumed to be constant, which causes the aircraft to have a constant shallow climb for a fixed angle of attack and Mach number.

The Breguet range equation is often used to quantify the performance of an aircraft. The initial and final weight of the aircraft is usually fixed for an aerodynamic design problem. Therefore, we can reduce the quantity of interests to $aM/c L/D$.

The contribute of the engine fuel consumption can be separated by assuming a constant thrust-specific fuel consumption for a range of Mach numbers, further reducing the quantity of interests to ML/D . In aircraft design, the 99% value of the maximum ML/D contour is often used to examine the robustness of the design [37]. The point with the highest Mach number on that contour line corresponds to the

long range cruise (LRC) point, which is the point at which the aircraft can fly at a higher speed by incurring a 1% increase in fuel burn [127].

However, the constant TSFC assumption is only accurate for a turbojet engine. The TSFC of a high bypass ratio turbofan varies with the square root of Mach number [126]. Therefore, if we substitute c with \sqrt{M} in the range equation, we can simplify the quantity of interests to $\sqrt{M}L/D$. In this appendix, we compare ML/D , $\sqrt{M}L/D$, and $aM/c L/D$ using a single-point optimized wing from Chapter 7. We add 100 drag counts to the computed drag to account for the drag due to the fuselage, tail, and nacelles.

B.2 Comparison of ML/D , $\sqrt{M}L/D$, and $aM/c L/D$ for a Single-Point Optimized Wing

We use the results of the single-point optimized wing design from Chapter 7 to perform the comparison. We plot the contours of each quantity of interests for a range of lift coefficients and Mach numbers, as shown in Figure B.1 to B.4.




Figure B.1: The contour of ML/D

The TSFC is computed based on a genetic GE90 engine model in Numerical Propulsion System Simulation (NPSS). We perform an inverse solve to compute the




Figure B.2: The contour of $\sqrt{M}L/D$




Figure B.3: The contour of $aM/c L/D$




Figure B.4: Comparison of the maxima locations and 99% contours

throttle setting for each flight conditions to match the thrust and drag coefficient. The blank area in Figure B.3 indicates the thrust required exceeded the limitation of the engine.

As shown in the figures, all three quantities of interests have similar trends and contour shapes. The contour and the maximum location of the $\sqrt{ML/D}$ is much closer to those of the $aM/c L/D$. The maximum of ML/D occurs at a higher Mach number and slightly lower C_L .

To exam the comparison at the operating condition, we take a 2-dimensional slice in the $M-C_L$ space for a constant $C_L = 0.5$. The results are shown in Figure B.5. Since the numerical values of each quantity are different, we normalize the lines based on the maximum values of each quantity. We see that both $\sqrt{ML/D}$ and $aM/c L/D$ have the same maximum location at Mach 0.850, which is the same as the Mach number of the single-point optimization in Chapter 7. The maximum of ML/D occurs at a slightly higher Mach number of 0.851. We conclude that all three quantities of interests show similar trends in Mach- C_L design space. However, $\sqrt{ML/D}$ should be used for aircraft with high bypass ratio turbofan to accurately predict the performance.




Figure B.5: $\sqrt{M}L/D$ and $aM/c L/D$ predicts the same maximum peak for $C_L = 0.5$.

APPENDIX C

Additional Data from Morphing Trailing Edge Optimization Results




Figure C.1: Morphing trailing edge optimization at $M = 0.85$, Alt = 28,000 ft, and $W = 347,500$ kg.




Figure C.2: Morphing trailing edge optimization at $M = 0.86$, Alt = 25,000 ft, and $W = 347,500$ kg.




Figure C.3: Morphing trailing edge optimization at $M = 0.86$, Alt = 33,000 ft, and $W = 347,500$ kg.




Figure C.4: Morphing trailing edge optimization at $M = 0.85$, Alt = 34,000 ft, and $W = 273,200$ kg.




Figure C.5: Morphing trailing edge optimization at $M = 0.85$, Alt = 28,000 ft, and $W = 273,200$ kg.




Figure C.6: Morphing trailing edge optimization at $M = 0.75$, Alt = 25,000 ft, and $W = 273,200$ kg.




Figure C.7: Morphing trailing edge optimization at $M = 0.75$, Alt = 37,000 ft, and $W = 273,200$ kg.




Figure C.8: Morphing trailing edge optimization at $M = 0.86$, Alt = 25,000 ft, and $W = 273,200$ kg.




Figure C.9: Morphing trailing edge optimization at $M = 0.86$, Alt = 37,000 ft, and $W = 273,200$ kg.




Figure C.10: Morphing trailing edge optimization at $M = 0.85$, $Alt = 39,000$ ft, and $W = 198,900$ kg.




Figure C.11: Morphing trailing edge optimization at $M = 0.75$, $Alt = 25,000$ ft, and $W = 198,900$ kg.




Figure C.12: Morphing trailing edge optimization at $M = 0.75$, Alt = 41,000 ft, and $W = 198,900$ kg.

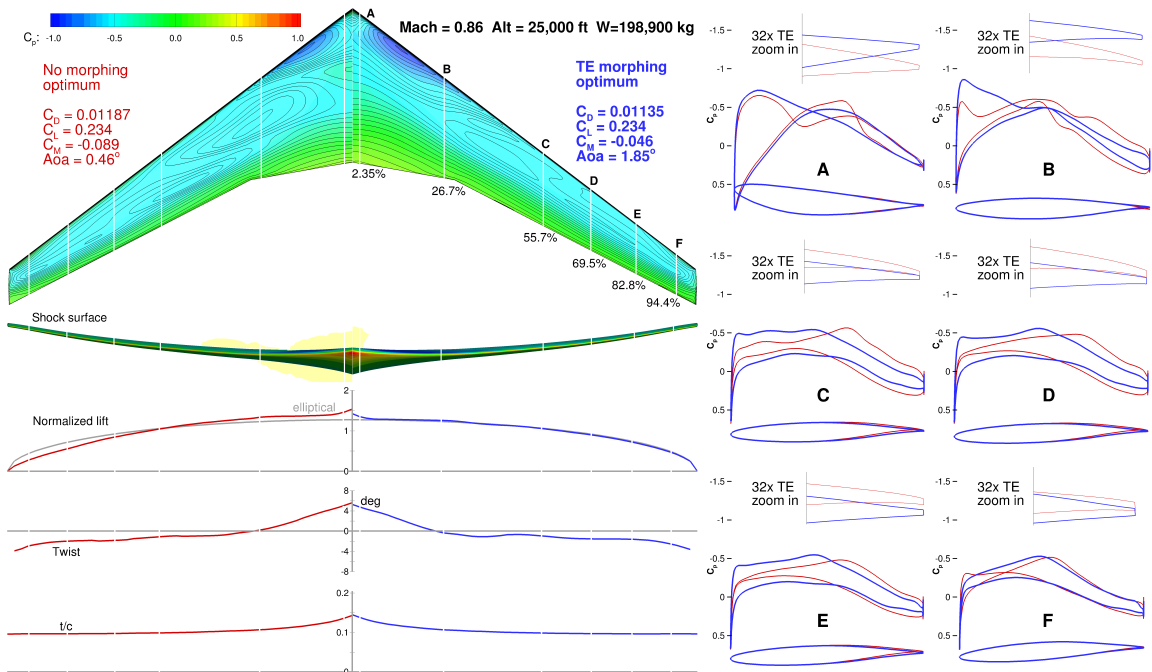


Figure C.13: Morphing trailing edge optimization at $M = 0.86$, Alt = 25,000 ft, and $W = 198,900$ kg.




Figure C.14: Morphing trailing edge optimization at $M = 0.86$, $Alt = 41,000$ ft, and $W = 198,900$ kg.




Figure C.15: Drag reduction contour for $W = 273,200$ kg with morphing TE.




Figure C.16: Drag reduction contour for $W = 273,200 \text{ kg}$ with full morphing wing.




Figure C.17: Drag reduction contour for $W = 198,900 \text{ kg}$ with morphing TE.




Figure C.18: Drag reduction contour for $W = 198,900$ kg with full morphing wing.




Figure C.19: ML/D contour for $W = 273,200$ kg of multipoint optimized baseline wing.




Figure C.20: ML/D contour for $W = 273,200$ kg of the morphing TE wing.




Figure C.21: ML/D contour for $W = 273,200$ kg of the full morphing wing.




Figure C.22: ML/D contour for $W = 198,900$ kg of multipoint optimized baseline wing.




Figure C.23: ML/D contour for $W = 198,900$ kg of the morphing TE wing.




Figure C.24: ML/D contour for $W = 198,900$ kg of the full morphing wing.

APPENDIX D

Euler-based Aerodynamic Design Optimization of a Blended-Wing-Body Aircraft

In this appendix, we present the results of Euler-based aerodynamic design optimization of the blended-wing-body aircraft.

D.1 Problem Formulation

Aerodynamic shape optimization of the BWB needs to be carefully formulated and constrained in order to achieve a physically feasible design. The following sections describe the objective function, design variables, and constraints.

D.1.1 Objective Function

We choose drag coefficient as the objective function for the optimization under prescribed lift. The drag coefficient has two components: drag coefficient from the Euler solver, and the skin friction drag coefficient. The van Driest method is used to capture the missing skin friction drag from the Euler solver. Skin friction drag is particularly important for trade-off between wing span and wing area.

D.1.2 Design Variables

Primary design variables are the geometric shape variables distributed on the FFD volume. A total of 800 shape variables are scattered on the lower and upper surfaces of the FFD volume, as shown in Figure [D.1](#). The large number of shape variables provides more degrees of freedom for the optimizer to explore, and to fine-tune the sectional airfoil shape and thickness-to-chord ratio at each spanwise location. We use

significantly more shape variables than in the previous studies in the literature. This is made possible by the implementation of an adjoint gradient calculation. We found that the cost of computing shape gradients is nearly independent of the number of shape variables.




Figure D.1: The BWB mesh (left), FFD volume and 800 shape control points (right)

The next set of design variables is the spanwise twist distribution. A total of 5 section twist design variables is used. The center of twist rotation is fixed at the reference axis, which is located at the quarter chord of each section. The twist variables provide a convenient way for the optimizer to minimize induced drag by adjusting the spanwise lift distribution, as well as to meet the root bending moment constraint.

Planform variables, such as span and sweep, are also considered in the optimization, and they contribute primarily to the reduction of induced drag and wave drag, respectively. The span design variable stretches the FFD volume in the spanwise direction. The sweep variable shears the FFD volume in chordwise direction. The planform variables can only be added to the optimization together with the bending moment constraint and the skin friction estimation, since the planform variation changes the structural weight and surface area significantly.

Auxiliary design variables are added to facilitate the formulation of the optimization problem. The angle-of-attack variable ensures that the lift coefficient constraint can be satisfied. We use an individual design feasible (IDF) approach to update the reference CG location (x_{CG}) and the mean aerodynamic chord (MAC). This requires the addition of a target variables, x_{CG}^t and MAC^t . Table D.1 summarizes the design variables.

Design Variables	Count	Design Variables	Count
Airfoil Shape	800	Angle-of-attack	1
Twist	5	Target CG location	1
Span	1	Target MAC	1
Sweep	1	Total	810

Table D.1: Summary of the design variables used in the BWB aerodynamic shape optimization problem

D.1.3 Constraints

As optimizers tend to explore any weakness in the numerical models, an optimization problem needs to be carefully constrained in order to yield a physically feasible design. Several geometric constraints are implemented. We impose thickness constraints near the leading edge (LE), trailing edge (TE), mid-chord, and centerbody to prevent the airfoil thickness from affecting low speed aerodynamic performance, to get a reasonable structural box depth, and to prevent the violation of the manufacturing constraints. These constraints also ensure sufficient height in the centerbody cabin and sufficient thickness at the LE and TE for the installation of high-lift devices such as slats and flaps. The volume of the centerbody is also constrained to meet the requirements for cabin and payload space. The wing volume constraints are also imposed to ensure sufficient space for fuel. In order to avoid generating non-physical kinked LE and TE, the shape variables located at the LE and TE are constrained so that each pair of shape variables can move only in opposite direction with equal magnitudes.

Due to the absence of a structural model, we use root bending moment as a surrogate for the structural weight trade-offs. The root bending moment is constrained to be equal or less than the baseline bending moment. With this constraint imposed, the optimized spanwise lift distribution tends less outboard loading instead of the elliptical distribution. The bending constraint is necessary to capture the trade-offs between aerodynamic performance and structural weight.

In addition, the BWB has to be trimmed in cruise conditions without the need to deflect its control surfaces which would result in trim drag. Therefore, a trim constraint is added. There are several ways to trim a flying wing: by unloading wing tips on a swept wing, by adding reflex to the airfoils at the trailing edge, and by adding anhedral to wing tips. Our optimization problem has all the required degrees of freedom to explore the design space except for the anhedral wing tips. In addition, stability is a particularly important design aspect of the BWB configuration. With

the absence of a conventional empennage, it is not immediately obvious whether a positive static margin can be achieved on a BWB aircraft. Therefore, we constrained the static margin to be greater than 5%. The static margin can be calculated as the ratio of the moment and lift derivatives.

$$K_n = -\frac{C_{M_\alpha}}{C_{L_\alpha}}, \quad (\text{D.1})$$

We calculate C_{M_α} and C_{L_α} using finite differences with an angle-of-attack step size of 0.1 degree. The static margin constraint incurs an additional computational cost. For each iteration, one additional flow solution and two additional adjoint solutions are required. Both flow and adjoint solutions have to be converged more accurately than usual to obtain an accurate static margin gradient. This is particularly important for static margin gradients with respect to shape variables, since they have relatively small magnitudes compared to other gradients.

The IDF formulation requires two additional compatibility constraints for the CG location and MAC. Table D.2 summarizes the constraints for the optimization problem.

Constraints	Count	Constraints	Count
Lift coefficient	1	Trim	1
Thickness	99	Static margin	1
Internal volume	99	CG compatibility	1
LE, TE control points	40	MAC compatibility	1
Total	243		

Table D.2: Summary of the constraints used in the BWB aerodynamic shape optimization problem

D.2 Aerodynamic Shape and Planform Optimization Results

The following sections discuss the results from performing a sequence of aerodynamic shape and planform optimizations. The baseline BWB geometry is used as the initial design. The optimizations are performed with an Euler solver and a friction drag estimation, as described in Section D.1. All cases are run on high performance computing clusters in the Center for Advanced Computing at the University of Michigan. Each computing node comprises two six-core 2.67 GHz Intel Xeon X5650 processors with 4 GB of RAM per core. All the nodes are interconnected with InfiniBand networking. A total of four cases are presented in this section.

- Case 1: Baseline lift-constrained drag minimization with respect to shape variables and with geometric constraints
- Case 2: Lift-constrained drag minimization with respect to shape variables, with geometric, root bending moment, and trim constraints
- Case 3: Lift-constrained drag minimization with respect to shape and planform variables, with geometric, root bending moment, and trim constraints
- Case 4: Lift-constrained drag minimization with respect to shape and planform variables, with geometric, root bending moment, trim, and static margin constraints

D.2.1 Case 1: baseline lift-constrained drag minimization with respect to shape variables and with geometric constraints

In order to validate our optimization formulation, we begin with a baseline optimization case with only geometric and lift constraints. The design variables are shape variables, twist, and angle-of-attack. Without any additional constraints, we expect to see an optimum elliptical lift distribution and weakened shocks. Since only one flow solution and two adjoint systems need to be solved, the problem requires less computational time than the other cases. The optimization is performed using 16 cores, and is converged in 34 hours with a total of 91 major optimization iterations. Figure D.2 shows the pressure distribution, airfoil shape, and lift, thickness-to-chord ratio (t/c), twist distributions of the baseline and optimized BWB configurations. The aerodynamic coefficients of the baseline and optimized BWB are listed in Table D.3. Figure D.3 shows the convergence history of the feasibility, optimality, and merit function.

Coefficient	Baseline	Optimized	Difference
C_D	0.02696	0.02245	-16.7%
C_L	0.440	0.440	0.0%
$C_{M_{root}}$	0.2599	0.2731	+5.1%
$C_{M_{trim}}$	0.1938	0.2037	+5.1%

Table D.3: Case 1: aerodynamic coefficients of the baseline and optimized BWB

The optimized BWB achieves an optimal elliptical spanwise lift distribution, mostly by varying the sectional twist distribution. Both root and tip of the wing are twisted down and the mid-span section is twisted up to match the elliptical lift.

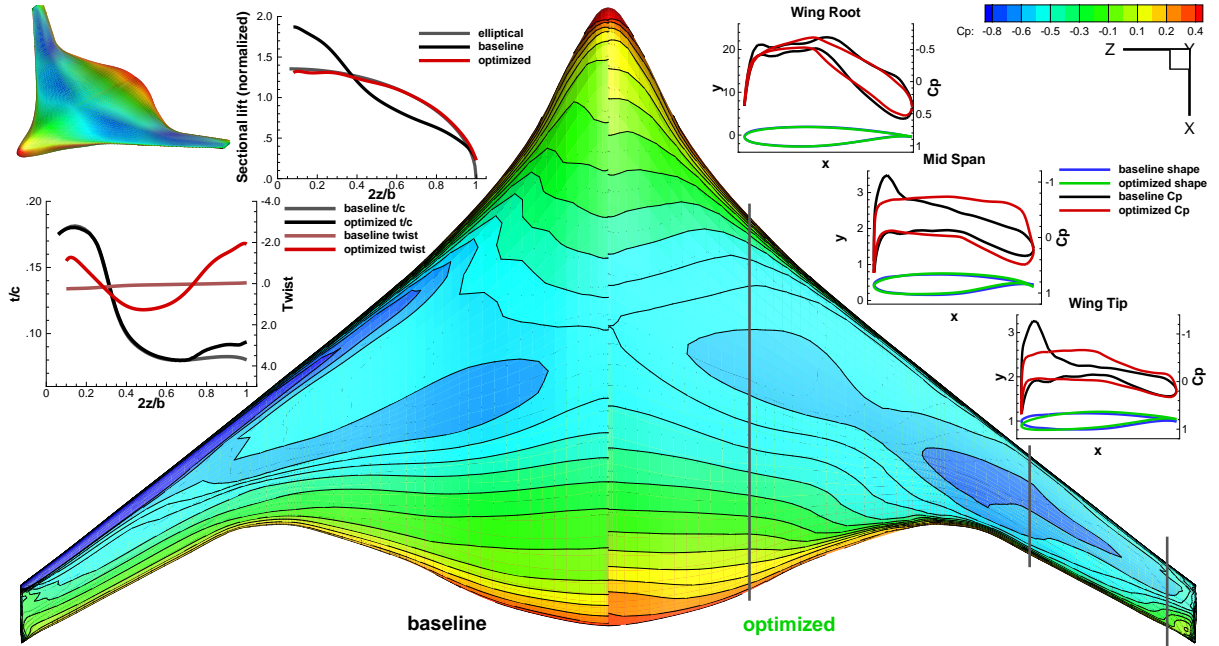


Figure D.2: Case 1: the results of the baseline lift constrained drag minimization with shape variables and geometric constraints. Pressure contour, sectional C_p distribution, airfoil shape, and spanwise lift, t/c , twist distributions of the baseline design (left) and the optimized design (right) are shown.




Figure D.3: Case 1: SNOPT optimization convergence history of the feasibility, optimality, and merit function

The airfoil shape of the centerbody is not significantly altered. Changes made to the shape of the centerbody are mainly to smooth out the pressure distribution. A large portion of lift is generated on the aft section of the centerbody because of the airfoil

camber, resulting in a pitch down moment. The optimized BWB further increases the pitch down moment by 5.1%, indicating the need for a trim constraint.

A strong shock can be seen at the leading edge of the upper wing section on the baseline BWB. The shock transits into a compression wave toward the centerbody. The shock structure on the optimized BWB is significantly different. The pressure distribution on the upper surface is flattened compared to that of the baseline BWB. The strong shock at the leading edge is weakened and shifted to the mid-chord, reducing the wave drag. The drag coefficient is reduced by 45 counts or 16.7%. With the optimal lift distribution, lift is shifted toward the wing tip. We can see that the root bending moment increases by 5.1%, which indicates that a heavier structure would be required. Therefore, a root bending moment constraint is necessary to constrain the impact on structural weight.

This baseline optimization serves as a validation of the optimization formulation. We obtain an elliptical lift distribution and reduced shock strength on the upper surface. We can also conclude that additional constraints are needed to achieve a practical design.

D.2.2 Case 2: lift-constrained drag minimization with respect to shape variables, with geometric, root bending moment, and trim constraints

As discussed in the previous case, we need to capture the trade-offs between the aerodynamic performance, structural weight, and trim drag. Therefore, root bending moment and trim constraints are added. The root bending moment is constrained to be less than or equal to the initial bending moment. Two additional adjoint systems are needed to compute the gradient of those constraints. This optimization problem is converged in 72 hours using 16 cores, and requires a total of 321 iterations. Figure D.4 shows the pressure distribution, airfoil shape, and lift, t/c, twist distributions of the baseline and optimized BWB configurations. Note that for the purpose of visualize the airfoil shape, x and y axes of the sectional airfoil plots are not of the same scale. Therefore, the sectional twist in these plots is not to scale. The aerodynamic coefficients of the baseline and optimized BWB are listed in Table D.4. Figure D.5 shows the convergence history of the feasibility, optimality, and merit function.

With the addition of root bending moment constraint, the optimized spanwise lift distribution is no longer elliptical. The lift on the wing section decreases more linearly toward the wing tip. This departure from the elliptical lift distribution limits span efficiency and hinders improvements in induced drag. However, it benefits the

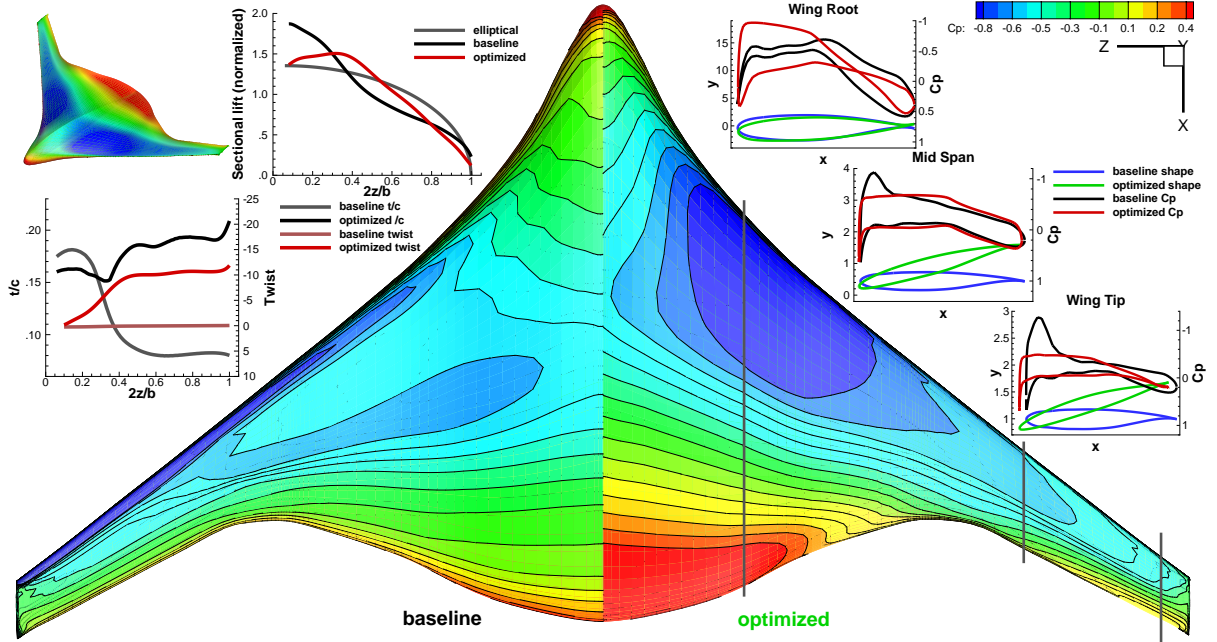


Figure D.4: Case 2: the results of the lift constrained drag minimization with shape variables, geometric, root bending moment, and trim constraints. Pressure contour, sectional C_p distribution, airfoil shape, and spanwise lift, t/c , twist distributions of the baseline design (left) and the optimized design (right) are shown.

Coefficient	Baseline	Optimized	Difference
C_D	0.02696	0.02381	-11.7%
C_L	0.440	0.440	0.0%
$C_{M_{root}}$	0.2599	0.2599	0.0%
$C_{M_{trim}}$	0.1938	0.0	-100%

Table D.4: Case 2: aerodynamic coefficients of the baseline and optimized BWB

structural weight and lateral control response for a flying wing. Two design features that lead to satisfaction of trim constraints on a trimmed flying wing are observed on the optimized BWB. The first is a reflex near the trailing edge of the optimized centerbody airfoil, resulting in a significant change to the chordwise lift distribution on the centerbody. All lift is generated at the forward section of the centerbody, while the aft centerbody has negative lift to trim the aircraft. As a result, the net lift on the centerbody is reduced, as shown on the lift distribution plot. The second feature is the unloaded wing tip. The optimized wing tip airfoil has washout and less lift than the baseline. The unloaded wing tips on a highly swept wing act as the horizontal tail to trim the aircraft.




Figure D.5: Case 2: SNOPT optimization convergence history of the feasibility, optimality, and merit function

Due to these constraints, optimized drag is reduced by 31 counts as compared to 45 counts in the previous case. This change is primarily due to a lower span efficiency and the presence of the reflex centerbody airfoil, which causes the lift induced drag and the trim drag to increase.

D.2.3 Case 3: lift-constrained drag minimization with respect to shape and planform variables, with geometric, root bending moment, and trim constraints

In this case, we seek the benefits of adding planform design variables to the optimization described in Case 2. By allowing changes to the span and sweep, we provide the optimizer with greater degrees of freedom and an efficient way of satisfying the trim constraints and reducing wave drag. This optimization problem is converged in 95 hours with 16 cores, and requires with a total of 426 iterations. The increased optimization time is mainly due to the increase in design space dimensionality. Figure D.6 shows the pressure distribution, airfoil shape, and lift, t/c , twist distributions of the baseline and optimized BWB. The aerodynamic coefficients of the baseline and optimized BWB are listed in Table D.5. Figure D.7 shows the convergence history of the feasibility, optimality, and merit function.

The sectional pressure distribution and pressure contour are relatively similar to those in the previous cases. The addition of planform variables does not significantly alter the pressure distribution. The optimized BWB also has a reflex centerbody

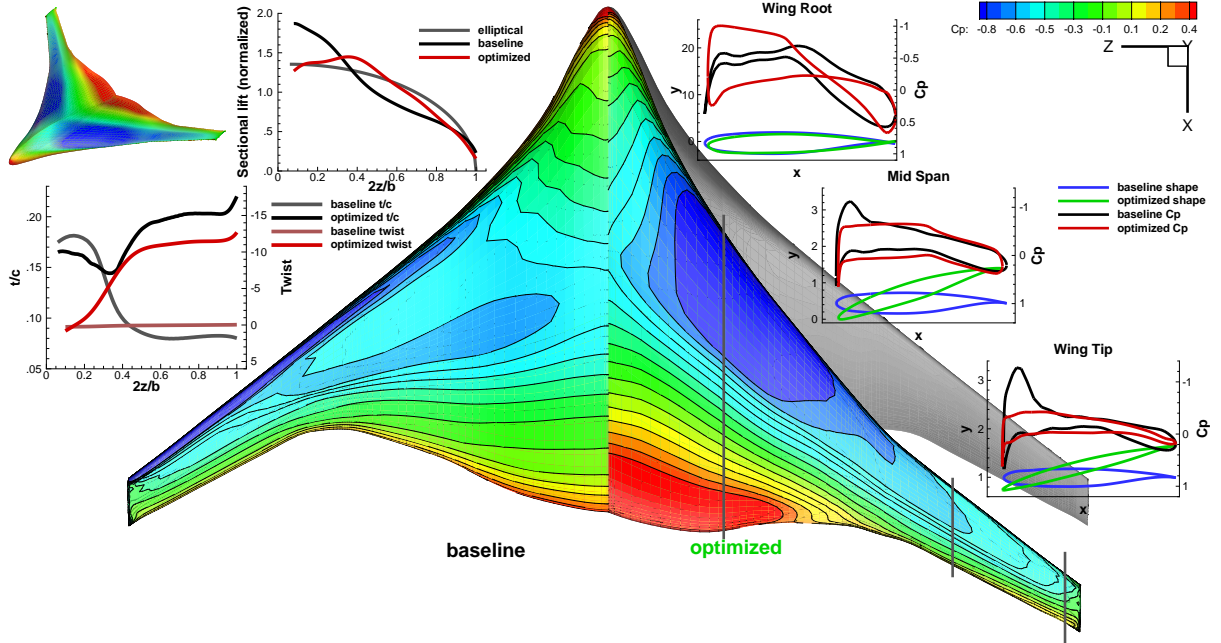


Figure D.6: Case 3: the results of the lift constrained drag minimization with shape and planform variables, geometric, root bending moment, and trim constraints. Pressure contour, sectional C_p distribution, airfoil shape, and spanwise lift, t/c , twist distributions of the baseline design (left) and the optimized design (right) are shown.

Coefficient	Baseline	Optimized	Difference
C_D	0.02696	0.02338	-13.3%
C_L	0.440	0.440	0.0%
$C_{M_{root}}$	0.2599	0.2599	0.0%
$C_{M_{trim}}$	0.1938	0.0	-100%

Table D.5: Case 3: aerodynamic coefficients of the baseline and optimized BWB

and an unloaded wing tip to trim the aircraft. The span of the optimized BWB is approximately the same as the baseline, while the sweep angle of the wing increased significantly. The increase in sweep angle reduces wave drag and also helps to trim the aircraft by increasing the moment arm of the unloaded tips. The wing tip generates a higher amount of lift compared to the previous case. The amount of centerbody reflex is reduced. The optimizer finds that trim with sweep and twist is more effective, which results in lower trim and induced drag. Increase in span is limited because of the root bending moment constraint and skin friction drag penalty.




Figure D.7: Case 3: SNOPT optimization convergence history of the feasibility, optimality, and merit function

D.2.4 Case 4: lift-constrained drag minimization with respect to shape and planform variables, with geometric, root bending moment, trim, and static margin constraints

The optimization setup for this case is similar to that of Case 3, with an exception of an additional static margin constraint. The goal of this case is to use numerical optimization to achieve an optimized longitudinally stable and trimmed configuration that would be otherwise difficult to design with a trial-and-error design process. The static margin is computed by performing finite difference of the lift and moment coefficients. We also added the payload location as a design variable, and this is represented by the CG location. The CG is allowed to move between 20% to 50% of the MAC.

This optimization problem is more computationally intensive than the previous cases for several reasons. Each iteration requires two flow solutions and six adjoint solutions in order to obtain the static margin and its gradient. The static margin gradient is a second order derivative, since it is the gradient of lift and moment coefficient gradients. Therefore, in order to achieve an accurate static margin gradient, both flow and adjoint solutions need to be converged to a higher tolerance (10^{-10}). Finite-differencing the lift and moment coefficients perturbs only the angle-of-attack. If a Newton-type iteration, such as the Newton–Krylov method, is used to solve the flow solution, the Maratos effect may cause the Newton iteration to stall. We increase number of Runge–Kutta iteration before switching to the Newton–Krylov solver. A

flow field solution reset is used as a fail-safe procedure if the solution stalls.

This optimization problem converged in 58 hours on 64 cores with a total of 138 iterations. Figure D.8 shows the CG and neutral point (NP) locations, pressure distribution, airfoil shape, and lift, t/c , twist distributions of the baseline and optimized BWB configurations. The aerodynamic coefficients of the baseline and optimized BWB are listed in Table D.6. Figure D.9 shows the convergence history of the feasibility, optimality, and merit function.

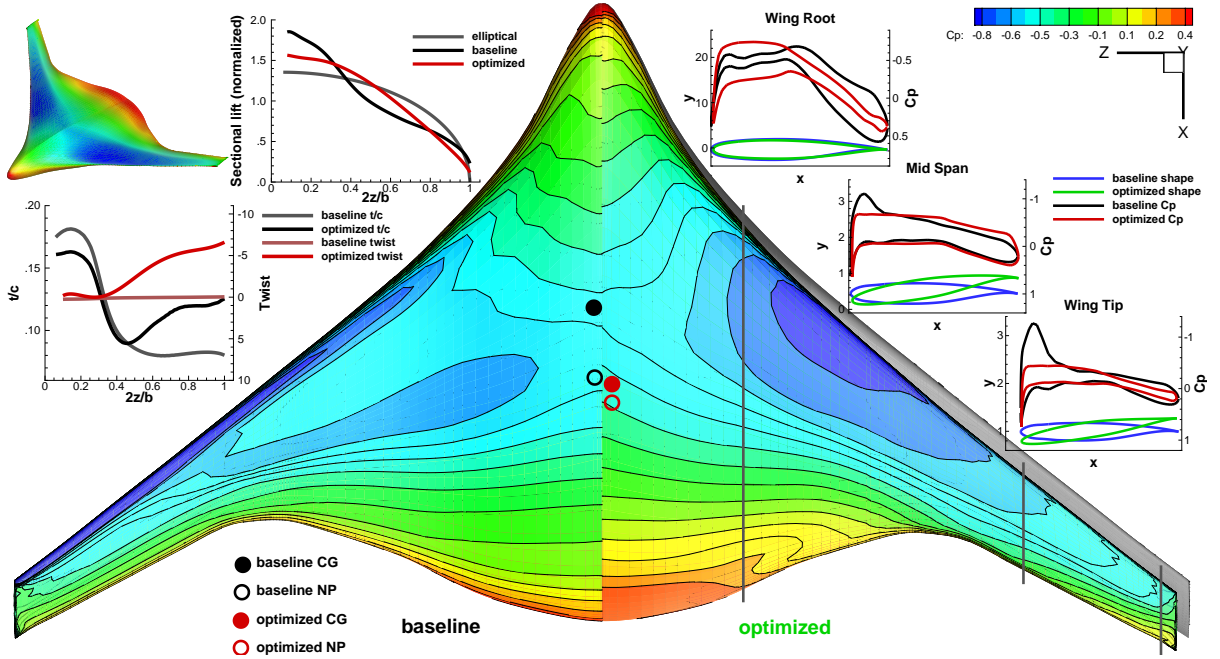


Figure D.8: Case 4: the results of the lift constrained drag minimization with shape and planform variables, geometric, root bending moment, trim, and static margin constraints. CG and NP locations, pressure contour, sectional C_p distribution, airfoil shape, and spanwise lift, t/c , twist distributions of the baseline design (left) and the optimized design (right) are shown.

Coefficient	Baseline	Optimized	Difference
C_D	0.02696	0.02339	-13.2%
C_L	0.440	0.440	0.0%
$C_{M_{root}}$	0.2599	0.2599	0.0%
$C_{M_{trim}}$	0.1938	0.0	-100%
K_n	19.1%	5.0%	/

Table D.6: Case 4: aerodynamic coefficients of the baseline and optimized BWB




Figure D.9: Case 4: SNOPT optimization convergence history of the feasibility, optimality, and merit function

Both static margin and trim constraints are sensitive to the CG location. Since a lower trim moment is required for an aft CG location, the trim constraint tends to move the CG backward by increasing CG design variable or by increasing the sweep angle. On the contrary, the CG location is also constrained by the static margin, since an aft CG location decreases the static margin. For a flying wing, location of the NP coincides with the aerodynamic center, the point about which the pitching moment coefficient does not vary with angle-of-attack. For a fixed planform, the aerodynamic center varies only slightly with the airfoil shape variables. The static margin constraint is most sensitive to the sweep angle. As sweep angle increases, both aerodynamic center and mean aerodynamic chord shift aft. However, the mean aerodynamic chord moves at a faster rate than the aerodynamic center. Thus, the static margin decreases for an increasing sweep. This effect can be seen from the optimization results. The sweep angle only increases slightly as compared to the previous case, which is constrained by the static margin. The static margin of the BWB is reduced from 19.1% in the baseline configuration to 5.0% in the optimized configuration. The CG moved from 25% of the MAC to 47%.

Because of the aft CG location than the previous cases, the trim constraint becomes less difficult to satisfy. We can see from the pressure distribution that the amount of reflex and down twist is reduced. Only aft centerbody has a reflex shape airfoil. The optimized BWB has better aerodynamic performance due to the relaxed CG location. It has a drag that is 36 counts lower than the baseline.

D.2.5 Summary of the Aerodynamic Shape and Planform Optimization Results

The four cases of aerodynamic shape optimization demonstrates the benefits and impacts of various design variables and constraints. As one would expect, the optimized BWB has the lowest drag coefficient when only lift and geometric constraints are enforced. However, this case also has a higher root bending moment and pitch moment than the baseline BWB. By adding trim and root bending moment constraints, the optimized drag coefficient increases by 14 counts relative to Case 1. The impact on the aerodynamic performance is mainly due to a less ideal spanwise lift distribution and increased drag from the reflex centerbody airfoil. We further explore the design space by adding span and sweep design variables to the optimization. The additional degrees-of-freedom provide about 4 drag count of improvement. This reduction comes from a lower skin friction drag and wave drag. Finally, we add both static margin constraint and CG design variable to the problem. We relax the CG location to allow the optimized BWB to meet the static margin constraint without a large impact on the aerodynamic performance. The optimized configuration has less sweep than the previous case, which is limited by the static margin constraint. Allowing the CG location to move improves the performance: the CG is shifted aft to reduce the reflex on the centerbody and the downwash on the wing tip. The optimized BWB maintains the same drag coefficient as the case without static margin. Table [D.7](#) summarizes the four aerodynamic shape optimization cases.

C_D	Reduction	Design Variables	Constraints				
			Geometric	Lift	Bending moment	Trim	Static margin
0.02696		Baseline	•	•			
0.02245	16.7%	Shape	•	•	•		
0.02381	11.7%	Shape	•	•	•	•	
0.02338	13.3%	Shape, Planform	•	•	•	•	
0.02339	13.2%	Shape, Planform, CG	•	•	•	•	•

Table D.7: Summary of optimal drag coefficients, design variables, and constraints of the aerodynamic shape optimization cases

BIBLIOGRAPHY

BIBLIOGRAPHY

- [1] McLean, D., *Understanding Aerodynamics: Arguing from the Real Physics*, Aerospace Series, Wiley, 2012.
- [2] Jameson, A., “Aerodynamic Design via Control Theory,” Vol. 3, No. 3, 1988, pp. 233–260. doi:[10.1007/BF01061285](https://doi.org/10.1007/BF01061285).
- [3] Reuther, J. J., Jameson, A., Alonso, J. J., Rimlinger, M. J., and Saunders, D., “Constrained Multipoint Aerodynamic Shape Optimization Using an Adjoint Formulation and Parallel Computers, Part 1,” *Journal of Aircraft*, Vol. 36, No. 1, 1999, pp. 51–60. doi:[10.2514/2.2413](https://doi.org/10.2514/2.2413).
- [4] Reuther, J. J., Jameson, A., Alonso, J. J., Rimlinger, M. J., and Saunders, D., “Constrained Multipoint Aerodynamic Shape Optimization Using an Adjoint Formulation and Parallel Computers, Part 2,” *Journal of Aircraft*, Vol. 36, No. 1, 1999, pp. 61–74. doi:[10.2514/2.2414](https://doi.org/10.2514/2.2414).
- [5] Leoviriyakit, K. and Jameson, A., “Multi-Point Wing Planform Optimization via Control Theory,” *43rd AIAA Aerospace Sciences Meeting and Exhibit*, 2005. doi:[10.2514/6.2005-450](https://doi.org/10.2514/6.2005-450).
- [6] Martins, J. R. R. A., Sturdza, P., and Alonso, J. J., “The Complex-Step Derivative Approximation,” *ACM Transactions on Mathematical Software*, Vol. 29, No. 3, 2003, pp. 245–262. doi:[10.1145/838250.838251](https://doi.org/10.1145/838250.838251).
- [7] Nadarajah, S. K., *The Discrete Adjoint Approach to Aerodynamic Shape Optimization*, Ph.D. thesis, Stanford University, 2003.
- [8] Amoignon, O., *Adjoint-based Aerodynamic Shape Optimization*, Ph.D. thesis, Uppsala University, 2003.
- [9] Nemec, M., Zingg, D. W., and Pulliam, T. H., “Multipoint and Multi-Objective Aerodynamic Shape Optimization,” *AIAA journal*, Vol. 42, No. 6, 2004, pp. 1057–1065. doi:[10.2514/1.10415](https://doi.org/10.2514/1.10415).
- [10] Hicken, J. E. and Zingg, D. W., “Aerodynamic Optimization Algorithm with Integrated Geometry Parameterization and Mesh Movement,” *AIAA journal*, Vol. 48, No. 2, 2010, pp. 400–413. doi:[10.2514/1.44033](https://doi.org/10.2514/1.44033).

- [11] Hicken, J. E. and Zingg, D. W., “Induced-Drag Minimization of Nonplanar Geometries Based on the Euler Equations,” *AIAA journal*, Vol. 48, No. 11, 2010, pp. 2564–2575. doi:[10.2514/1.J050379](https://doi.org/10.2514/1.J050379).
- [12] Osusky, L. and Zingg, D., “A Novel Aerodynamic Shape Optimization Approach for Three-Dimensional Turbulent Flows,” *50th AIAA Aerospace Sciences Meeting including the New Horizons Forum and Aerospace Exposition*, 2012. doi:[10.2514/6.2012-58](https://doi.org/10.2514/6.2012-58).
- [13] Lyu, Z. and Martins, J. R. R. A., “Aerodynamic Shape Optimization of a Blended-Wing-Body Aircraft,” *51st AIAA Aerospace Sciences Meeting including the New Horizons Forum and Aerospace Exposition*, 2013. doi:[10.2514/6.2013-283](https://doi.org/10.2514/6.2013-283).
- [14] Lyu, Z. and Martins, J. R. R. A., “RANS-based Aerodynamic Shape Optimization of a Blended-Wing-Body Aircraft,” *43rd AIAA Fluid Dynamics Conference and Exhibit*, June 2013.
- [15] Nadarajah, S. and Jameson, A., “Studies of the Continuous and Discrete Adjoint Approaches to Viscous Automatic Aerodynamic Shape Optimization,” *Proceedings of the 15th AIAA Computational Fluid Dynamics Conference*, Anaheim, CA, June 2001. doi:[10.2514/6.2001-2530](https://doi.org/10.2514/6.2001-2530).
- [16] Bischof, C., Corliss, G., Green, L., Griewank, A., Haigler, K., and Newman, P., “Automatic Differentiation of Advanced CFD Codes for Multidisciplinary Design,” *Computing Systems in Engineering*, Vol. 3, No. 6, 12 1992, pp. 625–637. doi:[10.1016/0956-0521\(92\)90014-A](https://doi.org/10.1016/0956-0521(92)90014-A).
- [17] Sherman, L. L., Taylor III, A. C., Green, L. L., Newman, P. A., Hou, G. W., and Korivi, V. M., “First-and Second-order Aerodynamic Sensitivity Derivatives via Automatic Differentiation with Incremental Iterative Methods,” *Journal of Computational Physics*, Vol. 129, No. 2, 1996, pp. 307–331. doi:[10.1006/jcph.1996.0252](https://doi.org/10.1006/jcph.1996.0252).
- [18] Mader, C. A., Martins, J. R. R. A., Alonso, J. J., and van der Weide, E., “ADjoint: An Approach for the Rapid Development of Discrete Adjoint Solvers,” *AIAA Journal*, Vol. 46, No. 4, April 2008, pp. 863–873. doi:[10.2514/1.29123](https://doi.org/10.2514/1.29123).
- [19] Spalart, P. and Allmaras, S., “A One-Equation Turbulence Model for Aerodynamic Flows,” *30th Aerospace Sciences Meeting and Exhibit*, 1992. doi:[10.2514/6.1992-439](https://doi.org/10.2514/6.1992-439).
- [20] Newton, S. I., *Philosophiæ Naturalis Principia Mathematica*, Londini, jussi Societatus Regiæ ac typis Josephi Streater; prostat apud plures bibliopolas, 1686.
- [21] Hicks, R. M., Murman, E. M., and Vanderplaats, G. N., “An Assessment of Airfoil Design by Numerical Optimization,” Tech. Rep. NASA-TM-X-3092, NASA, 1974.

- [22] Hicks, R. M. and Henne, P. A., “Wing Design by Numerical Optimization,” *Journal of Aircraft*, Vol. 15, 1978, pp. 407–412. doi:[10.2514/3.58379](https://doi.org/10.2514/3.58379).
- [23] Peter, J. E. V. and Dwight, R. P., “Numerical Sensitivity Analysis for Aerodynamic Optimization: A Survey of Approaches,” *Computers and Fluids*, Vol. 39, 2010, pp. 373–391. doi:[10.1016/j.compfluid.2009.09.013](https://doi.org/10.1016/j.compfluid.2009.09.013).
- [24] Martins, J. R. R. A. and Hwang, J. T., “Review and Unification of Methods for Computing Derivatives of Multidisciplinary Computational Models,” *AIAA Journal*, Vol. 51, No. 11, 2013, pp. 2582–2599. doi:[10.2514/1.J052184](https://doi.org/10.2514/1.J052184).
- [25] Pironneau, O., “On Optimum Profiles in Stokes Flow,” *Journal of Fluid Mechanics*, Vol. 59, No. 01, 1973, pp. 117–128. doi:[10.1017/S002211207300145X](https://doi.org/10.1017/S002211207300145X).
- [26] Pironneau, O., “On Optimum Design in Fluid Mechanics,” *Journal of Fluid Mechanics*, Vol. 64, 1974, pp. 97–110. doi:[10.1017/S0022112074002023](https://doi.org/10.1017/S0022112074002023).
- [27] Hazra, S., Schulz, V., Brezillon, J., and Gauger, N., “Aerodynamic Shape Optimization Using Simultaneous Pseudo-Timestepping,” *Journal of Computational Physics*, Vol. 204, No. 1, 2005, pp. 46 – 64. doi:[10.1016/j.jcp.2004.10.007](https://doi.org/10.1016/j.jcp.2004.10.007).
- [28] Lyu, Z., Kenway, G., Paige, C., and Martins, J. R. R. A., “Automatic Differentiation Adjoint of the Reynolds-Averaged Navier-Stokes Equations with a Turbulence Model,” *43rd AIAA Fluid Dynamics Conference and Exhibit*, June 2013.
- [29] Jameson, A., Martinelli, L., and Pierce, N., “Optimum Aerodynamic Design using the Navier–Stokes Equations,” *Theoretical and computational fluid dynamics*, Vol. 10, No. 1-4, 1998, pp. 213–237. doi:[10.1007/s001620050060](https://doi.org/10.1007/s001620050060).
- [30] Anderson, W. K. and Bonhaus, D. L., “Airfoil Design on Unstructured Grids for Turbulent Flows,” *AIAA journal*, Vol. 37, No. 2, 1999, pp. 185–191. doi:[10.2514/2.712](https://doi.org/10.2514/2.712).
- [31] Nielsen, E. J. and Anderson, W. K., “Aerodynamic Design Optimization on Unstructured Meshes Using the Navier-Stokes Equations,” *AIAA journal*, Vol. 37, No. 11, 1999, pp. 1411–1419. doi:[10.2514/2.640](https://doi.org/10.2514/2.640).
- [32] Dwight, R. P. and Brezillon, J., “Efficient and Robust Algorithms for Solution of the Adjoint Compressible Navier–Stokes Equations with Applications,” *International Journal for Numerical Methods in Fluids*, Vol. 60, 2009, pp. 365–389. doi:[10.1002/fld.1894](https://doi.org/10.1002/fld.1894).
- [33] Brezillon, J. and Dwight, R. P., “Applications of a Discrete Viscous Adjoint Method for Aerodynamic Shape Optimisation of 3D Configurations,” *CEAS Aeronautical Journal*, Vol. 3, No. 1, 2012, pp. 25–34. doi:[10.1007/s13272-011-0038-0](https://doi.org/10.1007/s13272-011-0038-0).

- [34] Nadarajah, S., Jameson, A., and Alonso, J. J., “An Adjoint Method for the Calculation of Remote Sensitivities in Supersonic Flow,” *Proceedings of the 40th Aerospace Sciences Meeting and Exhibit*, Reno, NV, January 2002. doi:[10.2514/6.2002-261](https://doi.org/10.2514/6.2002-261).
- [35] Nadarajah, S., Kim, S., Jameson, A., and Alonso, J., “Sonic Boom Reduction Using an Adjoint Method for Supersonic Transport Aircraft Configurations,” Vol. 73, 2003, pp. 355–362. doi:[10.1007/978-94-010-0017-8_53](https://doi.org/10.1007/978-94-010-0017-8_53).
- [36] Vassberg, J., “Introduction: Drag Prediction Workshop,” *Journal of Aircraft*, Vol. 45, No. 3, Jun 2008, pp. 737–737. doi:[10.2514/1.37761](https://doi.org/10.2514/1.37761).
- [37] Vassberg, J., Dehaan, M., Rivers, M., and Wahls, R., “Development of a Common Research Model for Applied CFD Validation Studies,” *26th AIAA Applied Aerodynamics Conference*, American Institute of Aeronautics and Astronautics, August 2008. doi:[10.2514/6.2008-6919](https://doi.org/10.2514/6.2008-6919).
- [38] Vassberg, J., “A Unified Baseline Grid about the Common Research Model Wing/Body for the Fifth AIAA CFD Drag Prediction Workshop (Invited),” *29th AIAA Applied Aerodynamics Conference*, Jul 2011. doi:[10.2514/6.2011-3508](https://doi.org/10.2514/6.2011-3508).
- [39] Kota, S., Hetrick, J., Osborn, R., Paul, D., Pendleton, E., Flick, P., and Tilmann, C., “Design and Application of Compliant Mechanisms for Morphing Aircraft Structures,” *Proceedings of SPIE*, Vol. 5054, 2003, p. 25.
- [40] Kota, S., Osborn, R., Ervin, G., Maric, D., Flick, P., and Paul, D., “Mission Adaptive Compliant Wing—Design, Fabrication and Flight Test,” *RTO Applied Vehicle Technology Panel (AVT) Symposium*, 2009.
- [41] Sofla, A. Y. N., Meguid, S. A., Tan, K. T., and Yeo, W. K., “Shape Morphing of Aircraft Wing: Status and Challenges,” *Materials & Design*, Vol. 31, No. 3, 3 2010, pp. 1284–1292. doi:[10.1016/j.matdes.2009.09.011](https://doi.org/10.1016/j.matdes.2009.09.011).
- [42] Molinari, G., Quack, M., Dmitriev, V., Morari, M., Jenny, P., and Ermanni, P., “Aero-Structural Optimization of Morphing Airfoils for Adaptive Wings,” *Journal of Intelligent Material Systems and Structures*, Vol. 22, No. 10, 2011, pp. 1075–1089. doi:[10.1177/1045389X11414089](https://doi.org/10.1177/1045389X11414089).
- [43] Lee, D., Gonzalez, L. F., Periaux, J., and Onate, E., “Robust Aerodynamic Design Optimisation of Morphing Aerofoil/Wing using Distributed MOGA,” *Proceedings of the 28th Congress of the International Council of the Aeronautical Sciences*, Optimage Ltd, on behalf of the International Council of the Aeronautical Sciences (ICAS)., 2012.
- [44] Liebeck, R. H., “Design of the Blended Wing Body Subsonic Transport,” *Journal of Aircraft*, Vol. 41, No. 1, May 2004, pp. 10–25. doi:[10.2514/1.9084](https://doi.org/10.2514/1.9084).

- [45] Kroo, I., “Innovations in Aeronautics,” *42nd AIAA Aerospace Sciences Meeting*, January 2004. doi:[0.2514/6.2004-1](https://doi.org/10.2514/6.2004-1).
- [46] Qin, N., Vavalle, A., Le Moigne, A., Laban, M., Hackett, K., and Weierfelt, P., “Aerodynamic Considerations of Blended Wing Body Aircraft,” *Progress in Aerospace Sciences*, Vol. 40, No. 6, 8 2004, pp. 321–343. doi:[10.1016/j.paerosci.2004.08.001](https://doi.org/10.1016/j.paerosci.2004.08.001).
- [47] Roman, D., Gilmore, R., and Wakayama, S., “Aerodynamics of High-Subsonic Blended-Wing-Body Configurations,” *41st Aerospace Sciences Meeting and Exhibit*, 2003. doi:[10.2514/6.2003-554](https://doi.org/10.2514/6.2003-554).
- [48] Liebeck, R., “Design of the Blended-Wing-Body Subsonic Transport,” *40th AIAA Aerospace Sciences Meeting & Exhibit*, 2002. doi:[10.2514/6.2002-2](https://doi.org/10.2514/6.2002-2).
- [49] Wakayama, S., “Multidisciplinary Design Optimization of the Blended-Wing-Body,” *7TH AIAA Symposium on Multidisciplinary Analysis And Optimization*, 1998. doi:[10.2514/6.1998-4938](https://doi.org/10.2514/6.1998-4938).
- [50] Wakayama, S., “Blended-Wing-Body Optimization Problem Setup,” *8th Symposium on Multidisciplinary Analysis and Optimization*, 2000. doi:[10.2514/6.2000-4740](https://doi.org/10.2514/6.2000-4740).
- [51] Qin, N., Vavalle, A., and Moigne, A. L., “Spanwise Lift Distribution for Blended Wing Body Aircraft,” *Journal of Aircraft*, Vol. 42, No. 2, 2005, pp. 356–365. doi:[10.2514/1.4229](https://doi.org/10.2514/1.4229).
- [52] Peigin, S. and Epstein, B., “Computational Fluid Dynamics Driven Optimization Of Blended Wing Body Aircraft,” *AIAA Journal*, Vol. 44, No. 11, 2006, pp. 2736–2745. doi:[10.2514/1.19757](https://doi.org/10.2514/1.19757).
- [53] Kuntawala, N. B., *Aerodynamic Shape Optimization of a Blended-Wing-Body Aircraft Configuration*, Master’s thesis, University of Toronto, 2011.
- [54] Kuntawala, N. B., Hicken, J. E., and Zingg, D. W., “Preliminary Aerodynamic Shape Optimization Of A Blended-Wing-Body Aircraft Configuration,” *49th AIAA Aerospace Sciences Meeting, Orlando, Florida*, 2011. doi:[10.2514/6.2011-642](https://doi.org/10.2514/6.2011-642).
- [55] Meheut, M., Arntz, A., and Carrier, G., “Aerodynamic Shape Optimizations of a Blended Wing Body Configuration for Several Wing Planforms,” *30th AIAA Applied Aerodynamics Conference*, 2012. doi:[10.2514/6.2012-3122](https://doi.org/10.2514/6.2012-3122).
- [56] Mader, C. A. and Martins, J. R. R. A., “Stability-Constrained Aerodynamic Shape Optimization of Flying Wings,” *Journal of Aircraft*, Vol. 50, No. 5, September 2013, pp. 1431–1449. doi:[10.2514/1.C031956](https://doi.org/10.2514/1.C031956).

- [57] Reist, T. A. and Zingg, D. W., “Aerodynamic Shape Optimization of a Blended-Wing-Body Regional Transport for a Short Range Mission,” *31st AIAA Applied Aerodynamics Conference*, Jul 2013. doi:[10.2514/6.2013-2414](https://doi.org/10.2514/6.2013-2414).
- [58] Kenway, G. K. W., Kennedy, G. J., and Martins, J. R. R. A., “Scalable Parallel Approach for High-Fidelity Steady-State Aeroelastic Analysis and Adjoint Derivative Computations,” *AIAA Journal*, Vol. 52, No. 5, 2014, pp. 935–951. doi:[10.2514/1.J052255](https://doi.org/10.2514/1.J052255).
- [59] Kenway, G. K. W. and Martins, J. R. R. A., “Multipoint High-fidelity Aerostructural Optimization of a Transport Aircraft Configuration,” *Journal of Aircraft*, Vol. 51, No. 1, 2014, pp. 144–160. doi:[10.2514/1.C032150](https://doi.org/10.2514/1.C032150).
- [60] Kennedy, G. J. and Martins, J. R. R. A., “A parallel aerostructural optimization framework for aircraft design studies,” *Structural and Multidisciplinary Optimization*, 2014. doi:[10.1007/s00158-014-1108-9](https://doi.org/10.1007/s00158-014-1108-9), (In press).
- [61] Samareh, J. A., “Survey of Shape Parameterization Techniques for High-Fidelity Multidisciplinary Shape Optimization,” *AIAA journal*, Vol. 39, No. 5, 2001, pp. 877–884. doi:[10.2514/2.1391](https://doi.org/10.2514/2.1391).
- [62] Kenway, G. K., Kennedy, G. J., and Martins, J. R. R. A., “A CAD-free Approach to High-Fidelity Aerostructural Optimization,” *Proceedings of the 13th AIAA/ISSMO Multidisciplinary Analysis Optimization Conference, Fort Worth, TX*, 2010. doi:[10.2514/6.2010-9231](https://doi.org/10.2514/6.2010-9231).
- [63] Thompson, J. F., Soni, B. K., and Weatherill, N. P., *Handbook of Grid Generation*, CRC press, 2010.
- [64] Batina, J. T., “Unsteady Euler Airfoil Solutions Using Unstructured Dynamic Meshes,” *AIAA journal*, Vol. 28, No. 8, 1990, pp. 1381–1388. doi:[10.2514/3.25229](https://doi.org/10.2514/3.25229).
- [65] van der Weide, E., Kalitzin, G., Schluter, J., and Alonso, J., “Unsteady Turbomachinery Computations Using Massively Parallel Platforms,” *44th AIAA Aerospace Sciences Meeting and Exhibit*, 2006. doi:[10.2514/6.2006-421](https://doi.org/10.2514/6.2006-421).
- [66] Jameson, A., Schmidt, W., and Turkel, E., “Numerical Solution of the Euler equations by Finite Volume Methods Using Runge Kutta Time Stepping Schemes,” *14th AIAA, Fluid and Plasma Dynamics Conference*, 1981.
- [67] Saad, Y. and Schultz, M. H., “GMRES: A Generalized Minimal Residual Algorithm for Solving Nonsymmetric Linear Systems,” *SIAM Journal on Scientific and Statistical Computing*, Vol. 7, No. 3, 1986, pp. 856–869. doi:[10.1137/0907058](https://doi.org/10.1137/0907058).
- [68] Balay, S., Gropp, W. D., McInnes, L. C., and Smith, B. F., “Efficient Management of Parallelism in Object Oriented Numerical Software Libraries,” *Modern*

Software Tools in Scientific Computing, edited by E. Arge, A. M. Bruaset, and H. P. Langtangen, Birkhäuser Press, 1997, pp. 163–202. doi:[10.1007/978-1-4612-1986-6_8](https://doi.org/10.1007/978-1-4612-1986-6_8).

- [69] Balay, S., Brown, J., , Buschelman, K., Eijkhout, V., Gropp, W. D., Kaushik, D., Knepley, M. G., McInnes, L. C., Smith, B. F., and Zhang, H., “PETSc Users Manual,” Tech. Rep. ANL-95/11 - Revision 3.4, Argonne National Laboratory, 2013.
- [70] Balay, S., Brown, J., Buschelman, K., Gropp, W. D., Kaushik, D., Knepley, M. G., McInnes, L. C., Smith, B. F., and Zhang, H., “PETSc Web page,” 2013, <http://www.mcs.anl.gov/petsc>.
- [71] Gill, P. E., Murray, W., and Saunders, M. A., “SNOPT: An SQP Algorithm for Large-Scale Constrained Optimization,” *SIAM journal on optimization*, Vol. 12, No. 4, 2002, pp. 979–1006. doi:[10.1137/S1052623499350013](https://doi.org/10.1137/S1052623499350013).
- [72] Perez, R. E., Jansen, P. W., and Martins, J. R. R. A., “pyOpt: A Python-Based Object-Oriented Framework for Nonlinear Constrained Optimization,” *Structures and Multidisciplinary Optimization*, Vol. 45, No. 1, 2012, pp. 101–118. doi:[10.1007/s00158-011-0666-3](https://doi.org/10.1007/s00158-011-0666-3).
- [73] Kraft, D., “A Software Package for Sequential Quadratic Programming,” Tech. rep., Tech. Rep. DFVLR-FB 88-28, DLR German Aerospace Center, 1988.
- [74] Lawson, C. L. and Hanson, R. J., *Solving Least Squares Problems*, Vol. 161, SIAM, 1974.
- [75] Kraft, D., “Algorithm 733: TOMP–Fortran Modules for Optimal Control Calculations,” *ACM Transactions on Mathematical Software (TOMS)*, Vol. 20, No. 3, 1994, pp. 262–281.
- [76] Wächter, A. and Biegler, L. T., “On the Implementation of an Interior-point Filter Line-search Algorithm for Large-scale Nonlinear programming,” *Mathematical programming*, Vol. 106, No. 1, 2006, pp. 25–57.
- [77] Vanderplaats, G. N., *CONMIN, a FORTRAN Program for Constrained Function Minimization: User’s Manual*, Vol. 62282, Ames Research Center and US Army Air Mobility R&D Laboratory, 1973.
- [78] Svanberg, K., “The Method of Moving Asymptotes—a New Method for Structural Optimization,” *International journal for numerical methods in engineering*, Vol. 24, No. 2, 1987, pp. 359–373.
- [79] Jansen, P. W. and Perez, R. E., “Constrained Structural Design Optimization via a Parallel Augmented Lagrangian Particle Swarm Optimization Approach,” *Computers & Structures*, Vol. 89, No. 13–14, 7 2011, pp. 1352–1366. doi:[10.1016/j.compstruc.2011.03.011](https://doi.org/10.1016/j.compstruc.2011.03.011).

- [80] Perez, R. and Behdinan, K., “Particle Swarm Approach for Structural Design Optimization,” *Computers & Structures*, Vol. 85, No. 19, 2007, pp. 1579–1588.
- [81] Jansen, P., Perez, R. E., and Martins, J. R. R. A., “Aerostructural Optimization of Nonplanar Lifting Surfaces,” *Journal of Aircraft*, Vol. 47, No. 5, 2010, pp. 1491–1503. doi:[10.2514/1.44727](https://doi.org/10.2514/1.44727).
- [82] Haghghat, S., Martins, J. R. R. A., and Liu, H. H. T., “Aeroservoelastic Design Optimization of a Flexible Wing,” *Journal of Aircraft*, Vol. 49, No. 2, 2012, pp. 432–443. doi:[10.2514/1.C031344](https://doi.org/10.2514/1.C031344).
- [83] Deb, K., Pratap, A., Agarwal, S., and Meyarivan, T., “A Fast and Elitist Multiobjective Genetic Algorithm: NSGA-II,” *Evolutionary Computation, IEEE Transactions on*, Vol. 6, No. 2, 2002, pp. 182–197. doi:[10.1109/4235.996017](https://doi.org/10.1109/4235.996017).
- [84] Rall, L. B. and Corliss, G. F., “An Introduction to Automatic Differentiation,” *Computational Differentiation: Techniques, Applications, and Tools*, 1996, pp. 1–17.
- [85] Bischof, C., Roh, L., and Mauer-Oats, A., “ADIC: an Extensible Automatic Differentiation Tool for ANSI-C,” *Software—Practice & Experience*, Vol. 27, No. 12, 1997, pp. 1427–1456. doi:[10.1002/\(SICI\)1097-024X\(199712\)27:12<1427::AID-SPE138>3.3.CO;2-H](https://doi.org/10.1002/(SICI)1097-024X(199712)27:12<1427::AID-SPE138>3.3.CO;2-H).
- [86] Bischof, C., Khademi, P., Mauer, A., and Carle, A., “Adifor 2.0: Automatic Differentiation of Fortran 77 Programs,” *Computational Science Engineering, IEEE*, Vol. 3, No. 3, 1996, pp. 18–32. doi:[10.1109/99.537089](https://doi.org/10.1109/99.537089).
- [87] Bendtsen, C. and Stauning, O., “FADBAD, a Flexible C++ Package for Automatic Differentiation,” Tech. rep., Department of Mathematical Modelling, Technical University of Denmark, 1996.
- [88] Utke, J., Naumann, U., Fagan, M., Tallent, N., Strout, M., Heimbach, P., Hill, C., and Wunsch, C., “OpenAD/F: A Modular Open-Source Tool for Automatic Differentiation of Fortran Codes,” *ACM Trans. Math. Softw.*, Vol. 34, No. 4, July 2008, pp. 18:1–18:36. doi:[10.1145/1377596.1377598](https://doi.org/10.1145/1377596.1377598).
- [89] Hascoët, L., “TAPENADE: a Tool for Automatic Differentiation of Programs,” *Proceedings of 4th European Congress on Computational Methods, ECCOMAS*, 2004, pp. 1–14.
- [90] Nielsen, E. J. and Kleb, W. L., “Efficient Construction of Discrete Adjoint Operators on Unstructured Grids using Complex Variables,” *AIAA journal*, Vol. 44, No. 4, 2006, pp. 827–836. doi:[10.2514/1.15830](https://doi.org/10.2514/1.15830).
- [91] Goldfarb, D. and Toint, P. L., “Optimal Estimation of Jacobian and Hessian Matrices that Arise in Finite Difference Calculations,” *Mathematics of Computation*, Vol. 43, No. 167, 1984, pp. 69–88. doi:[10.2307/2007400](https://doi.org/10.2307/2007400).

- [92] Schmitt, V. and Charpin, F., “Pressure Distributions on the ONERA-M6-Wing at Transonic Mach Numbers,” *Experimental Data Base for Computer Program Assessment*, 1979, pp. B1–1.
- [93] Obayashi, S. and Guruswamy, G. P., “Convergence Acceleration of a Navier-Stokes Solver for Efficient Static Aeroelastic Computations,” *AIAA Journal*, Vol. 33, No. 6, 2013/05/30 1995, pp. 1134–1141. doi:[10.2514/3.12533](https://doi.org/10.2514/3.12533).
- [94] Mani, M., Ladd, J., Cain, A., Bush, R., Mani, M., Ladd, J., Cain, A., and Bush, R., “An Assessment of One- and Two-Equation Turbulence Models for Internal and External Flows,” *28th Fluid Dynamics Conference*, 1997. doi:[10.2514/6.1997-2010](https://doi.org/10.2514/6.1997-2010).
- [95] Nielsen, E. J. and Anderson, W. K., “Recent Improvements in Aerodynamic Design Optimization on Unstructured Meshes,” *AIAA journal*, Vol. 40, No. 6, 2002, pp. 1155–1163. doi:[10.2514/2.1765](https://doi.org/10.2514/2.1765).
- [96] Moigne, A. L. and Qin, N., “Variable-Fidelity Aerodynamic Optimization for Turbulent Flows Using a Discrete Adjoint Formulation,” *AIAA Journal*, Vol. 42, No. 7, 2004, pp. 1281–1292. doi:[10.2514/1.2109](https://doi.org/10.2514/1.2109).
- [97] Rho, O., Lee, K., Kim, C., Kim, C., and Lee, B., “Parallelized design optimization for transonic wings using aerodynamic sensitivity analysis,” *40th AIAA Aerospace Sciences Meeting & Exhibit*, 2002. doi:[10.2514/6.2002-264](https://doi.org/10.2514/6.2002-264).
- [98] Osusky, M. and Zingg, D., “A parallel Newton-Krylov-Schur flow solver for the Reynolds-Averaged Navier-Stokes equations,” *50th AIAA Aerospace Sciences Meeting including the New Horizons Forum and Aerospace Exposition*, 2012. doi:[10.2514/6.2012-442](https://doi.org/10.2514/6.2012-442).
- [99] Bueno-Orovio, A., Castro, C., Palacios, F., and Zuazua, E., “Continuous Adjoint Approach for the Spalart-Allmaras Model in Aerodynamic Optimization,” *AIAA Journal*, Vol. 50, No. 3, 2012, pp. 631–646. doi:[10.2514/1.J051307](https://doi.org/10.2514/1.J051307).
- [100] Lyu, Z. and Martins, J. R. R. A., “Aerodynamic Design Optimization Studies of a Blended-Wing-Body Aircraft,” *Journal of Aircraft*, Vol. 51, No. 5, 2014, pp. 1604–1617. doi:[10.2514/1.C032491](https://doi.org/10.2514/1.C032491).
- [101] Lyu, Z., Kenway, G. K. W., and Martins, J. R. R. A., “Aerodynamic Shape Optimization Investigations of the Common Research Model Wing Benchmark,” *AIAA Journal*, 2014. doi:[10.2514/1.J053318](https://doi.org/10.2514/1.J053318), (In press).
- [102] Lyu, Z. and Martins, J. R. R. A., “Aerodynamic Shape Optimization of an Adaptive Morphing Trailing Edge Wing,” *15th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*, Atlanta, GA, June 2014. doi:[10.2514/6.2014-3275](https://doi.org/10.2514/6.2014-3275), AIAA 2014-3275.

- [103] Sasaki, D., Morikawa, M., Obayashi, S., and Nakahashi, K., “Aerodynamic Shape Optimization of Supersonic Wings by Adaptive Range Multiobjective Genetic Algorithms,” 2001, pp. 639–652. doi:[10.1007/3-540-44719-9_45](https://doi.org/10.1007/3-540-44719-9_45).
- [104] Foster, N. F. and Dulikravich, G. S., “Three-dimensional Aerodynamic Shape Optimization using Genetic and Gradient Search Algorithms,” *Journal of Spacecraft and Rockets*, Vol. 34, No. 1, 1997, pp. 36–42.
- [105] Zingg, D. W., Nemec, M., and Pulliam, T. H., “A Comparative Evaluation of Genetic and Gradient-Based Algorithms Applied to Aerodynamic Optimization,” *European Journal of Computational Mechanics*, Vol. 17, No. 1–2, January 2008, pp. 103–126. doi:[10.3166/remn.17.103-126](https://doi.org/10.3166/remn.17.103-126).
- [106] Obayashi, S. and Tsukahara, T., “Comparison of optimization algorithms for aerodynamic shape design,” *AIAA journal*, Vol. 35, No. 8, 1997, pp. 1413–1415.
- [107] Frank, P. D. and Shubin, G. R., “A Comparison of Optimization-based Approaches for a Model Computational Aerodynamics Design Problem,” *Journal of Computational Physics*, Vol. 98, No. 1, 1992, pp. 74–89.
- [108] Lyu, Z., Xu, Z., and Martins, J. R. R. A., “Benchmarking Optimization Algorithms for Wing Aerodynamic Design Optimization,” *8th International Conference on Computational Fluid Dynamics (ICCFD8)*, Chengdu, China, July 2014, ICCFD8-2014-0203.
- [109] Rosenbrock, H. H., “An Automatic Method for Finding the Greatest or Least Value of a Function,” *The Computer Journal*, Vol. 3, No. 3, 1960, pp. 175–184. doi:[10.1093/comjnl/3.3.175](https://doi.org/10.1093/comjnl/3.3.175).
- [110] Moré, J. J., Garbow, B. S., and Hillstom, K. E., “Testing Unconstrained Optimization Software,” *ACM Transactions on Mathematical Software (TOMS)*, Vol. 7, No. 1, 1981, pp. 17–41.
- [111] Hackbusch, W., *Multi-Grid Methods and Applications*, Vol. 4, Springer-Verlag Berlin, 1985.
- [112] Turkel, E., “Review of Preconditioning Methods for Fluid Dynamics,” *Applied Numerical Mathematics*, Vol. 12, No. 1–3, 1993, pp. 257 – 284. doi:[10.1016/0168-9274\(93\)90122-8](https://doi.org/10.1016/0168-9274(93)90122-8).
- [113] Gropp, W., Keyes, D., McInnes, L. C., and Tidriri, M., “Globalized Newton-Krylov-Schwarz Algorithms and Software for Parallel Implicit CFD,” *International Journal of High Performance Computing Applications*, Vol. 14, No. 2, 2000, pp. 102–136. doi:[10.1177/109434200001400202](https://doi.org/10.1177/109434200001400202).
- [114] Knoll, D. and Keyes, D., “Jacobian-free Newton–Krylov Methods: a Survey of Approaches and Applications,” *Journal of Computational Physics*, Vol. 193, No. 2, 2004, pp. 357 – 397. doi:[0.1016/j.jcp.2003.08.010](https://doi.org/10.1016/j.jcp.2003.08.010).

- [115] Lyu, Z., Kenway, G. K. W., and Martins, J. R. R. A., “RANS-based Aerodynamic Shape Optimization Investigations of the Common Research Model Wing,” *AIAA Science and Technology Forum and Exposition (SciTech)*, National Harbor, MD, January 2014.
- [116] Vassberg, J. and Jameson, A., “Influence of Shape Parameterization on Aerodynamic Shape Optimization,” Tech. rep., Von Karman Institute, Brussels, Belgium, April 2014.
- [117] Telidetzki, K., Osusky, L., and Zingg, D. W., “Application of Jetstream to a Suite of Aerodynamic Shape Optimization Problems,” *52nd Aerospace Sciences Meeting*, Feb 2014. doi:[10.2514/6.2014-0571](https://doi.org/10.2514/6.2014-0571).
- [118] Carrier, G., Destarac, D., Dumont, A., Meheut, M., Din, I. S. E., Peter, J., Khelil, S. B., Brezillon, J., and Pestana, M., “Gradient-Based Aerodynamic Optimization with the elsA Software,” *52nd Aerospace Sciences Meeting*, Feb 2014. doi:[10.2514/6.2014-0568](https://doi.org/10.2514/6.2014-0568).
- [119] Lyu, Z. and Martins, J. R. R. A., “Strategies for Solving High-Fidelity Aerodynamic Shape Optimization Problems,” *15th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*, Atlanta, GA, June 2014. doi:[10.2514/6.2014-2594](https://doi.org/10.2514/6.2014-2594), AIAA 2014-2594.
- [120] Roache, P. J., “Verification of Codes and Calculations,” *AIAA Journal*, Vol. 36, No. 5, 1998, pp. 696–702. doi:[10.2514/2.457](https://doi.org/10.2514/2.457).
- [121] Gill, P. E., Murray, W., and Saunders, M. A., *User’s Guide for SNOPT Version 7: Software for Large-Scale Nonlinear Programming*, 2006.
- [122] Holst, T. L. and Pulliam, T. H., “Aerodynamic Shape Optimization using a Real-Number-Encoded Genetic Algorithm,” *NASA Technical Reports*, 2001.
- [123] Chernukhin, O. and Zingg, D. W., “Multimodality and Global Optimization in Aerodynamic Design,” *AIAA Journal*, April 2013, pp. 1–13. doi:[10.2514/1.J051835](https://doi.org/10.2514/1.J051835).
- [124] Bisson, F. and Nadarajah, S., “Adjoint-Based Aerodynamic Design Optimization Using the Drag Decomposition Method,” *31st AIAA Applied Aerodynamics Conference*, July 2013. doi:[10.2514/6.2013-2909](https://doi.org/10.2514/6.2013-2909).
- [125] Liem, R. P., Kenway, G. K. W., and Martins, J. R. R. A., “Multi-mission Aircraft Fuel Burn Minimization via Multi-point Aerostructural Optimization,” *AIAA Journal*, 2014, (Accepted).
- [126] Torenbeek, E., *Advanced Aircraft Design: Conceptual Design, Technology and Optimization of Subsonic Civil Airplanes*, Aerospace Series, Wiley, 2013. doi:[10.1002/9781118568101](https://doi.org/10.1002/9781118568101).

- [127] Roberson, W., Root, R., and Adams, D., “Fuel Conservation Strategies: Cruise Flight,” Tech. rep., Boeing AERO Magazine, 2007.
- [128] Haimes, R., “Automated Feature Extraction from Transient CFD Simulations,” *Proceeding of the 7th Annual Conference of the CFD Society of Canada, Halifax, NS*, May 1999.
- [129] Martins, J. R. R. A. and Lambe, A. B., “Multidisciplinary Design Optimization: A Survey of Architectures,” *AIAA Journal*, Vol. 51, No. 9, Aug 2013, pp. 2049–2075. doi:[10.2514/1.J051895](https://doi.org/10.2514/1.J051895).
- [130] Mader, C. A. and Martins, J. R. R. A., “Computation of Aircraft Stability Derivatives Using an Automatic Differentiation Adjoint Approach,” *AIAA Journal*, Vol. 49, No. 12, 2011, pp. 2737–2750. doi:[10.2514/1.J051147](https://doi.org/10.2514/1.J051147).
- [131] Mader, C. A. and Martins, J. R. R. A., “Optimal Flying Wings: A Numerical Optimization Study,” *53rd AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference*, Honolulu, HI, April 2012. doi:[10.2514/6.2012-1758](https://doi.org/10.2514/6.2012-1758).
- [132] Jones, R. T., “The Spanwise Distribution of Lift for Minimum Induced Drag of Wings Having a Given Lift and A Given Bending Moment,” TN-2249, NASA, Dec. 1950.
- [133] Liem, R. P., Kenway, G. K., and Martins, J. R. R. A., “Multi-Point, Multi-Mission, High-Fidelity Aerostructural Optimization of a Long-Range Aircraft Configuration,” *Proceedings of the 14th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*, Indianapolis, IN, Sep 2012. doi:[10.2514/6.2012-5706](https://doi.org/10.2514/6.2012-5706).