

High-Fidelity Load and Gradient Corrections for Static Aeroelastic Tailoring of Composite Wings

Jovanov, Kristofer

DOI

[10.4233/uuid:14b55d5e-586a-4641-8990-55a397674db8](https://doi.org/10.4233/uuid:14b55d5e-586a-4641-8990-55a397674db8)

Publication date

2019

Document Version

Final published version

Citation (APA)

Jovanov, K. (2019). *High-Fidelity Load and Gradient Corrections for Static Aeroelastic Tailoring of Composite Wings*. [Dissertation (TU Delft), Delft University of Technology].
<https://doi.org/10.4233/uuid:14b55d5e-586a-4641-8990-55a397674db8>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

**High-Fidelity Load and Gradient
Corrections for Static Aeroelastic
Tailoring of Composite Wings**

High-Fidelity Load and Gradient Corrections for Static Aeroelastic Tailoring of Composite Wings

Dissertation

for the purpose of obtaining the degree of doctor
at Delft University of Technology
by the authority of the Rector Magnificus,
prof. dr. ir. T.H.J.J. van der Hagen,
chair of the Board for Doctorates
to be defended publicly on
Wednesday 3 July 2019 at 10:00 o'clock.

by

Kristofer JOVANOV

Master of Science in Engineering
Chalmers University of Technology, Sweden
born in Gothenburg, Sweden.

This dissertation has been approved by the promotors.

Composition of the doctoral committee:

Rector Magnificus,	chairperson
Prof. dr. C. Bisagni,	promotor Delft University of Technology
Dr. ir. R. De Breuker,	promotor Delft University of Technology

Independent members:

Prof. dr. M. Karpel	Technion - Israel Institute of Technology
Prof. dr. V.V. Toropov	Queen Mary University of London
Prof. dr. ir. H. Bijl	Leiden University
Prof. dr. G.J.W. van Bussel	Delft University of Technology

Reserve member:

Prof. dr. ir. R. Benedictus	Delft University of Technology
-----------------------------	--------------------------------

Other members:

Mr. C. Blondeau	ONERA, France
-----------------	---------------

This research work was supported by Delft University of Technology and the AMEDEO project, part of the European Union Seventh Framework Program and a Marie Curie Initial Training Network.



AMEDEO project partners:



Keywords: Optimization, Aeroelasticity, Sensitivity Analysis

Cover: Design by Kristofer Jovanov. CFD model is courtesy of ONERA.

Copyright © 2019 by Kristofer Jovanov

ISBN 978-94-028-1565-8

An electronic version of this dissertation is available at
<http://repository.tudelft.nl/>.

SUMMARY

Multi-disciplinary design optimization (MDO) is a field that has gained traction in recent years. It can be distinguished as a methodology that promotes a simultaneous change in design features associated to multiple disciplines in order to achieve the best possible design of a coupled multi-disciplinary problem. A typical application in aircraft design is aero-structural design optimization, where the interaction of fluids and structures on load-carrying components is considered. The recent progress in this field can be largely attributed to the increasing computational resources and parallel computing in the form of High Performance Computing Clusters. Nevertheless, the computational bottleneck which persists to this day in high-fidelity aero-structural design is the gradient computation, which is a fundamental component in the design of systems with several thousand design variables and constraints.

There are several methods that address the issue of computing coupled aero-structural gradients for large-scale problems. The most popular approach is to formulate functions of interest, which are typically element stresses or strains, as a single function by constraint aggregation. In doing so, the number of large linear systems that must be solved to obtain the required gradients can be reduced to the same amount as the number of constraint aggregation functions. This method however, albeit straightforward and efficient, can result in inaccurate gradients and by extension in overly conservative designs.

The aim of this thesis is not to develop a high-fidelity MDO framework, as there are already several known successful attempts. Such an endeavor would require much time and add little novelty compared to existing state-of-the-art. However, the gradient computation of the coupled aero-structural problem has been identified as an Achilles heel in many existing frameworks and the main effort in this doctoral thesis is devoted to improving this important facet in high-fidelity aero-structural design.

A framework is initially designed in the thesis to compute static aeroelastic analyses of wings in transonic conditions using the finite element method to model the structure and Euler flow aerodynamics to model the surrounding pressure. The convergence of the static aeroelastic problem is improved by a vortex lattice aerodynamic model using a defect-correction strategy. A similar method is developed to solve the subsequent gradient problem where the vortex lattice model

features as a preconditioner to accelerate the solution of the linear aero-structural gradient problem. The method is applied to structural sizing variables and only the direct method is considered. The proposed method is shown to outperform conventional methods such as linear block Gauss-Seidel with various types of relaxation. The method is successfully demonstrated on coupled aero-structural gradient computations with Euler and RANS flow aerodynamics.

Computing gradients by solving the linear aero-structural gradient problem one design variable at a time is not computationally viable when the design problem is subject to thousands of constraints and design variables. A method is developed that leverages information from the vortex lattice model to reconstruct the high-fidelity aero-structural gradients. Singular value decomposition in conjunction with pseudo-inverting is used to reconstruct the high-fidelity gradients, and it is shown that the reconstructed gradients are improved when the vector basis resulting from the vortex lattice model is enriched.

The proposed gradient reconstruction method is demonstrated on static aeroelastic tailoring of a forward-swept wing. The objective is weight minimization, the constraints are limited to strain and buckling, whereas lamination parameters and laminate thicknesses are employed as design variables. The gradient reconstruction method indicates that the design results in a lower structural weight when compared to an equivalent design procedure where the gradients are constructed purely by vortex lattice aerodynamics. Moreover, the enrichment of basis vectors in the gradient reconstruction scheme is shown to improve the convergence behavior of the optimization problem.

SAMENVATTING

Multidisciplinaire optimalisatie (MDO) is een vakgebied dat recentelijk steeds meer aandacht krijgt. Het is een methode die ontwerpvariabelen van verschillende disciplines combineert om aldus het beste ontwerp te creëren voor een multidisciplinair probleem. Een typische toepassing in vliegtuigontwerp is aero-structurele optimalisatie die de interactie beschrijft tussen de stroming en de belastingen op dragende onderdelen. Dankzij verbeterde hardware van en parallelle rekenmethoden op supercomputers is dit vakgebied met rasse schreden vooruitgegaan. Desondanks is de beperkende factor voor aero-structurele optimalisatie met hoge nauwkeurigheid de berekening van de afgeleiden. Deze afgeleiden zijn een essentieel onderdeel voor het ontwerpen van systemen met duizenden ontwerpvariabelen en randvoorwaarden.

Meerdere methoden zijn in staat om de gekoppelde aero-structurele afgeleiden van grote optimalisatieproblemen te berekenen. De meest populaire aanpak is om functies van randvoorwaarden samen te voegen tot één enkele functie door randvoorwaardencollectie. Op deze manier wordt het aantal lineaire systemen dat opgelost moet worden gereduceerd tot het aantal randvoorwaardencollectiefuncties. Hoewel dit een eenvoudige en efficiënte methode is, kan deze methode leiden tot onnauwkeurige afgeleiden en veel te conservatieve ontwerpen.

Het is niet het doel van deze thesis om een MDO-rekenmethode te ontwerpen met hoge nauwkeurigheid, want er bestaan al talrijke werkende methoden. Dit zou erg veel tijd kosten en weinig nieuws toevoegen aan de huidige stand van zaken. Echter, de Achilleshiel van menig rekenmethode is de berekening van de gekoppelde aero-structurele afgeleiden. Daarom is dan ook het hoofddoel van deze thesis om deze berekening efficiënter te maken.

Een methode is aanvankelijk ontwikkeld voor een statische aeroelastische analyse van vleugels in een transsonne stroming met behulp van de eindige elementenmethode voor de structurele berekeningen en Euler aerodynamica voor de stromingsberekeningen. Een defect-correctiemethode met behulp van een vortexmethode is aangewend om de convergentie van het statisch aeroelastische probleem te verbeteren. Een gelijkaardige methode is ontwikkeld voor de berekening van de bijbehorende afgeleiden. In dat geval wordt de vortexmethode gebruikt als conditionerend element om het oplossen van de lineaire aero-structurele afgeleiden te versnellen. Deze methode is toegepast op structurele ontwerpvariabelen en enkel

de directe methode wordt gebruikt. Het is aangetoond dat deze methode beter werkt dan conventionele methoden zoals daar zijn de lineaire block Gauss-Seidel methode met verschillende relaxatiemethoden. De methode is succesvol toegepast op gekoppelde aero-structurele afgeleideberekeningen met Euler en RANS stromingen.

Het is rekenkundig niet haalbaar om zulk lineair systeem op te lossen voor elke variabele apart omdat het hele ontwerpprobleem duizenden randvoorwaarden en ontwerpvariabelen telt. Een methode is ontwikkeld om afgeleiden met hoge nauwkeurigheid te reconstrueren met behulp van vortexmethoden. Dit is bewerkstelligd door het gebruik van de decompositie van de singuliere waarden samen met pseudo-inverteren om de afgeleiden met hoge nauwkeurigheid te reconstrueren en het is aangetoond dat de gereconstrueerde afgeleiden verbeteren naar gelang de vectorbasis, die gebaseerd is op vortexmethoden, verrijkt wordt.

De methode om afgeleiden te reconstrueren is toegepast op het statisch aeroelastisch ontwerpen van een voorwaarts gepijlde vleugel. Het doel van de optimalisatie is gewichtsbesparing en de randvoorwaarden zijn maximale rek en knikbelasting. De ontwerpvariabelen zijn de laminaatparameters en laminaatdiktes. De methode voor het reconstrueren van afgeleiden toont aan dat een lager gewicht bereikt kan worden met deze methode in vergelijking met het gebruik van afgeleiden die berekend zijn met vortexmethoden. Het verrijken van de vectorbasis verbetert aantoonbaar de convergentie van de optimalisatie.

NOMENCLATURE

ABBREVIATIONS

AIC	Aerodynamic Influence Coefficient
BPR	Bypass Ratio
BWE	Backward-Euler
CAD	Computer-Aided Design
CFD	Computational Fluid Dynamics
CRM	Common Research Model
DLM	Doublet Lattice Method
DOF	Degrees of Freedom
FEM	Finite Element Method/Model
FGMRES	Flexible Generalized Minimal RESidual
FSI	Fluid-Structure Interaction
FSW	Forward-Swept Wing
GCMMA	Globally Convergent Method of Moving Asymptotes
GMRES	Generalized Minimal RESidual
GSE	Global Sensitivity Equations
HPC	High Performance Computing
ILU	Incomplete LU-factorization
IPS	Infinite-Plate Spline
KS	Kreisselmeier-Steinhauser
LBGS	Linear Block Gauss-Seidel
LC	Load Case
LE	Leading Edge
LU-SSOR	Lower-Upper Symmetric Successive OverRelaxation
MAC	Mean Aerodynamic Chord
MDO	Multidisciplinary Design Optimization
MF	Multifidelity

MUSCL	Monotonic Upwind Scheme for Conservation Laws
NLBGS	NonLinear Block Gauss-Seidel
OM6	ONERA M6 (wing)
RANS	Reynolds-Averaged Navier-Stokes
RBF	Radial Basis Function
RPK	Revenue Passenger Kilometer
SVD	Singular Value Decomposition
TE	Trailing Edge
TPS	Thin-Plate Spline
VLM	Vortex Lattice Method

SYMBOLS

$[\mathbf{H}]$	Mesh deformation operator
$[\mathbf{T}]$	Load transfer operator
\mathcal{R}_a	Aerodynamic residual
\mathcal{R}_s	Structural residual
$\mathbf{\Gamma}$	Vortex strength vector
σ	Element stresses
\mathbf{A}_{IC}	Aerodynamic influence coefficient matrix
\mathbf{K}	Structural stiffness matrix
\mathbf{K}_a	Aerodynamic stiffness matrix
\mathbf{p}	Design variable vector
\mathbf{Q}/\mathbf{Q}_s	Structural load vector
\mathbf{Q}_a	Aerodynamic load vector
\mathbf{S}	Stress-displacement matrix
\mathbf{U}	Structural degrees of freedom (structural states)
\mathbf{V}_b	Flow tangency boundary condition vector
\mathbf{V}_∞	Free-stream velocity vector
\mathbf{V}_{tot}	Total velocity vector
\mathbf{W}	Conservative variables (aerodynamic states)
\mathbf{X}_a	Aerodynamic grid coordinates
\mathbf{X}_p	Panel grid coordinates
\mathbf{X}_s	Structural grid coordinates
\mathcal{J}	Function of interest
ρ_∞	Free-stream density

CONTENTS

SUMMARY	I
SAMENVATTING	III
NOMENCLATURE	V
1 INTRODUCTION	1
1.1 MULTI-DISCIPLINARY DESIGN OPTIMIZATION	5
1.2 SENSITIVITY ANALYSIS	7
1.3 RESEARCH OBJECTIVE	9
1.4 THESIS LAYOUT	11
2 GRADIENT COMPUTATION	13
2.1 COMPUTATIONAL MODULES	14
2.1.1 Aerodynamic solvers	14
2.1.2 Structural solver	15
2.1.3 Coupling module	17
2.1.4 VLM aerodynamic stiffness matrix	19
2.2 AEROELASTIC ANALYSIS	20
2.3 SENSITIVITY ANALYSIS	22
2.4 CASE STUDY	31
2.4.1 Setup	31
2.4.2 Results	32

2.5 CONCLUSIONS	39
3 GRADIENT RECONSTRUCTION	41
3.1 AEROELASTIC ANALYSIS AND SENSITIVITY ANALYSIS	42
3.2 GRADIENT RECONSTRUCTION	44
3.2.1 Methodology	44
3.2.2 SVD with respect to the energy norm	46
3.2.3 Pseudo-inverse with respect to the energy norm	48
3.2.4 Solution procedure	49
3.3 CASE STUDY I - ONERA M6	49
3.4 CASE STUDY II - FORWARD SWEEP WING	53
3.4.1 Extension to trim analysis and sensitivity analysis	55
3.4.2 Results	59
3.5 CONCLUSIONS	62
4 STATIC AEROELASTIC TAILORING OF A FORWARD-SWEEP WING	63
4.1 MODEL DESCRIPTION	65
4.2 OPTIMIZATION SETUP	68
4.2.1 Load cases	68
4.2.2 Objective, design variables and constraints	69
4.2.3 Optimizer and initial design	70
4.2.4 Optimization cases	71
4.3 OPTIMIZATION RESULTS	72
4.3.1 Airfoil influence	74
4.3.2 Euler load and gradient corrections	79
4.4 CONCLUSIONS	83
5 CONCLUSIONS AND RECOMMENDATIONS	85
5.1 CONCLUSIONS	86
5.2 RECOMMENDATIONS	88

BIBLIOGRAPHY

91

ACKNOWLEDGEMENTS

99

1

INTRODUCTION

Civil aviation has come a long way since its infancy in the 1940s. Considered a luxury in the early days, air travel has today become an indispensable means of transportation in modern societies. This trend is continuously expanding to emerging markets where civil aviation has not prospered to the same extent. Indeed, the two juggernauts of civil aviation, Boeing and Airbus, predict a continued growth in air travel in the coming years. In its "*Current Market Outlook 2017-2036*" [1], Boeing identifies the Asian market as the primary driver for the growing demand in air travel. According to its market forecast, the Asian jet fleet is set to expand nearly threefold from currently 6,830 units to 17,520 units in the next twenty years. Boeing's main competitor, Airbus, predicts similar trends in its "*Global Market Forecast 2017-2036*" [2]. More importantly, both companies envision an average annual Revenue Passenger Kilometer (RPK) growth of $\approx 5\%$ in the next twenty years.

The growing demand in air travel is likely to stimulate economies across the globe. However, as the air traffic intensifies, it becomes all the more important to curb global warming by limiting emissions of greenhouse gases such as carbon dioxides (CO_2), but also reducing air pollutants primarily in form of nitrogen oxides (NO_x). To this end, several target goals were established by the European Commission in agreement with key stakeholders from the aircraft industry in "*ACARE Flightpath 2050*" [3]. The primary goals are a 75% reduction in CO_2 and a 90% reduction in NO_x emissions per RPK by 2050 compared to new aircraft in 2000. The goals brought forward were indeed aspirational with the aim of pushing existing, but

also developing technologies, into reaching the target goals. To better understand the challenges that lie ahead in fulfilling the aforementioned goals, it is always beneficial to take a look back and reflect upon past achievements in civil aviation.

In a recent technical report, Kharina and Rutherford [4] analyzed the fuel efficiency of commercial jets during 1968-2014. The average fuel burn per passenger kilometer during the said period is depicted in Figure 1.1. There are some interesting trends worth highlighting. The sudden drop in the late 1960s, for instance, can be largely attributed to the introduction of Boeing 747. It was the most fuel efficient aircraft at its time with the first high Bypass Ratio (BPR) turbofan engines. However, sales of the B747 saturated over time as smaller, less fuel efficient jets gained ground. Introduced in 1976, the supersonic Concorde for instance, albeit considered an engineering marvel, had a fuel burn per passenger kilometer multiple times that of the B747. This, among other factors, caused the fuel burn to increase and then stagnate in the 1970s. A second period of stagnation can also be observed in the late 1990s, followed by a fuel burn reduction around 2004. This can be largely attributed to the soaring oil prices in the early 2000s. It became less affordable for airline companies to place new orders with increasing oil prices, which prompted manufacturers to develop more fuel efficient aircraft. In 2006, Airbus outlined plans for an improved version of their best selling A320 with a target efficiency gain of 4-5%. Boeing followed suit a few years later by announcing the decision to re-engineer the B737, with similar gains.

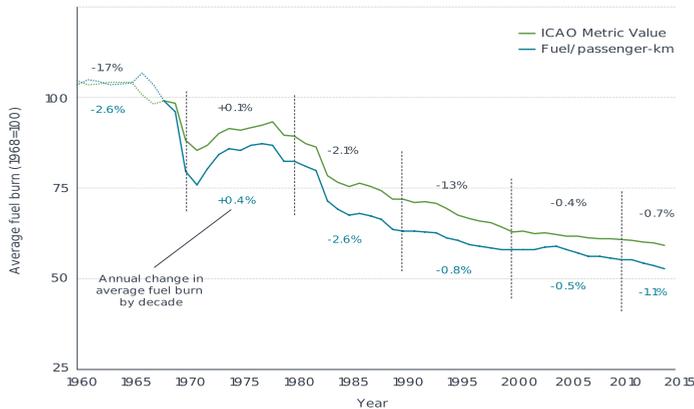


Figure 1.1: Average fuel burn for new commercial jet aircraft [4]

In spite of limited periods of marginal or no improvements, the overall trend is very encouraging with a 45% fuel burn reduction during 1968-2014. A significant contribution to this trend, in addition to the aforementioned high BPR turbofan engines, can be owed to the advancements in aerodynamic design. In the early days, aerodynamic design was predominantly exercised on a trial-and-error basis. This principle is neatly summed up by a quote from Samuel Beckett: *“Ever tried. Ever failed. No matter. Try again. Fail again. Fail better”*. A lot of time and ef-

fort was dedicated to building mockups, performing wind tunnel tests and gaining empirical knowledge. As Computational Fluid Dynamics (CFD) gained traction throughout the 1960s, an increasing amount of time was reallocated from the wind tunnel to the work station. Today, CFD has matured into an established tool in industry for preliminary design. Reynolds-Averaged Navier-Stokes (RANS) models that were once considered computationally expensive, are fast becoming the standard in aerodynamic shape optimization primarily due to increased computational resources. An overview of the evolution of CFD, its impact on the engineering design process and its future potential is provided by Jameson [5].

Another development which carries great potential of improving the fuel burn efficiency, by reducing the aircraft weight, is the growing prevalence of composite materials. Possessing a high strength-to-weight ratio, these materials can be integrated in structural optimization to reduce the weight of the airframe. Composite materials have only recently started being used on a larger scale with the introduction of the new flagship models Boeing B787 and Airbus A350. Figure 1.2 is a case in point of the dramatic change in material composition in favor of composites. The key to this change is the technological advancement in the production line where primary structures, such as the fuselage and the wings, are now manufactured by sophisticated tape laying machines.

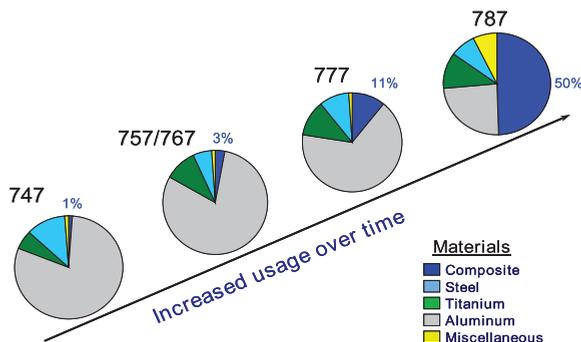


Figure 1.2: Increased use of composites over time [6]

As much as aerodynamic design and structural optimization have advanced in aircraft industry, they have predominantly done so separated from one another. Paul Rubbert [7], a former chief aerodynamicist at Boeing, argues that the development of cutting edge CFD tools in aerodynamic design has taken precedence over the ability to make proper cross-functional trades involving disciplines such as loads, structures, control, noise, manufacturing and so forth. This, in effect, has rendered the design process to become sequential. For instance, the structural skeleton of a wing is often sized only after an ideal aerodynamic shape has been determined. However, aerodynamics and structures are under specific circumstances strongly interconnected disciplines. Consider the wing of any commercial aircraft as an example. Aerodynamic pressure applied on the wing will result in

a structural deformation. As the wing deforms from its initial configuration, the pressure will no longer be valid but must be re-evaluated at the deformed state. With a new pressure field, the structure will deviate from its initial deformation and so forth. This is a classic example of Fluid-Structure Interaction (FSI) or as it appears to be more commonly referred to in the aeronautical community: *aero-structural analysis*. In a sequential design process, the aerodynamicist will seek to improve the aerodynamic performance by reducing the drag, typically by increasing the span so as to reduce the induced drag. However, this often contradicts the objective of the structural engineer who seeks to reduce the aircraft weight. An elongated wing must be reinforced in order to avoid structural failure, hence adding more weight. Consequently, there exists a trade-off between aerodynamic and structural performance which must be carefully considered.

In general, a computational model is only as good as its predictive capability of the real-life events it is trying to describe. If aero-structural coupling effects are considerable, as is the case with aircraft wings, and yet neglected by the model, then any optimized results based on said model are bound to be flawed or suboptimal at best. To achieve superior designs that can accurately reflect real physical behavior, three important criteria should be considered:

1. The computational models should be of "high-fidelity" character, i.e. they should capture all important physical behavior that might occur at a given flight condition. This is not to say that "low-fidelity" models are of less importance. Rather, the model in question should capture all the necessary physical features under the given conditions.
2. Any inter-disciplinary coupling effects, if present, should be properly accounted for by the computational model(s).
3. Design optimization including multiple disciplines should be performed simultaneously, as opposed to sequentially. A global multi-disciplinary objective should be constructed that applies a carefully considered trade-off on the individual single-disciplinary objectives in order to maximize the operating performance based on the flight mission.

The third criterion encapsulates the concept of Multidisciplinary Design Optimization (MDO). It entails the simultaneous improvement of multiple disciplines (e.g. aerodynamic performance and structural weight) by considering changes in the design that influence both disciplines. For instance, changes in the outer aerodynamic shape of the wing must influence the geometry of internal structural components such as ribs, spars and stiffeners, and vice versa. MDO has come to be widely regarded in recent years as the next stepping stone in the engineering design process. It has the potential to facilitate the development of novel, more fuel efficient concepts that deviate from the standard cylindrical fuselage configuration of today's commercial aircraft.

1.1 MULTI-DISCIPLINARY DESIGN OPTIMIZATION

An early effort to demonstrate the advantages of MDO was made by Grossman et al. [8]. In this work, a sailplane wing was modeled by lifting line aerodynamics coupled to a beam structural model. The design variables were limited to planform parameters and element thicknesses. Based on the said configuration, they successfully showed that the designs obtained by integrated optimization (i.e. MDO) were superior in terms of both aerodynamic performance and structural weight, compared to designs obtained by sequential optimization.

In a follow-up paper [9], they extended the modeling complexity to vortex lattice aerodynamics and implemented a more detailed finite element model including membrane, shear and rod elements. A forward-swept wing configuration was considered in this work. Moreover, the inter-disciplinary gradients were computed more efficiently with an improved optimization algorithm. The structural design variables considered not only element thicknesses, but also composite materials by a single ply orientation variable. The geometric design variables were similar to the previous work.

With increasing computational resources in the 1990s, high-fidelity aerodynamics started gaining traction. Lifting line theory or models based on linear potential flow theory were superseded by Euler and Navier-Stokes models. Structural beam models were replaced by complex wingbox structural models including varying finite elements types (e.g. shear, shell, bars and so forth). For example, Chattopadhyay and Pagaldipti [10] coupled a box beam structural model to a Navier-Stokes flow model, whereas Baker and Giesing [11] coupled a detailed finite element structural model to an Euler flow model. The work in both papers was aimed at investigating MDO of high-speed civil transport aircraft. The improvements in sensitivity analysis and the optimization algorithms allowed for computational models of higher complexity and computation cost to be employed. However, the number of design variables remained relatively low.

Giunta et al. [12] proposed an alternative approach to including high-fidelity models in MDO. In this work, the design space was restricted to the region of interest by numerical models of lower complexity. The MDO of the aircraft was subsequently performed by response surfaces at the points of interest using numerical models of higher complexity. The proposed approach significantly reduced the computational cost and mitigated noise present in the original models from the optimization process.

Maute et al. [13] designed an MDO framework based on Euler CFD and a linear finite element model including spars, ribs, hinges and control surfaces. The two disciplines were coupled in a three-field formulation, with the mesh deformation of the aerodynamic grid regarded as the third field. A simplified Computer-Aided Design (CAD) model was used to parametrize the surface geometry from

1

which the computational meshes were derived. The optimization considered sweep and twist changes as design variables, in addition to structural thicknesses of the wing. Different strategies for computing the analytical sensitivity derivatives were presented. It was shown that the strategy with an exact linearization of the flow problem in the sensitivity analysis outperformed the alternative strategies with an approximated flux Jacobian in the optimization of the Aeroelastic Research Wing. The additional computational cost of computing the exact gradients was compensated for by a reduction in the number of optimization iterations.

Reuther et al. [14] and Martins et al. [15] were amongst the first to implement adjoint-based aero-structural sensitivity analysis in a high-fidelity model framework consisting of Euler/RANS aerodynamics coupled to a finite element wingbox model. Ten twist and 180 airfoil shape variables were considered in a drag minimization problem for a three dimensional wing. In subsequent work [16], 76 shape variables associated to twist, camber and bump functions were considered in addition to 10 structural thickness variables for an optimization of a supersonic business jet. The MDO approach resulted in a 16% structural weight reduction of the aircraft compared to sequential optimization.

Martins and co-workers further improved their multi-disciplinary design optimization framework. Several additional computational improvements were presented by Kenway et al. [17]. The computational efficiency of the MDO framework was improved by monolithic solution algorithms and enhanced parallelism. The aim in their work was to optimize transport aircraft, considering multiple points in the flight envelope [18]. A detailed structural finite element model was coupled to an Euler flow solver with a free-form deformation method to parametrize the geometry. In recent work by Martins and co-workers, efforts have been made to extend the MDO framework to include composite materials [19], buckling constraints [20] and flutter constraints [21].

Much of the recent work on aero-structural optimization is centered on high-fidelity models (e.g. Euler or RANS CFD) while maintaining a feasible computational cost by implementing improved algorithms. However, it is also recognized that these high-fidelity optimization frameworks are still computationally expensive in preliminary design, and much more so in conceptual design, as multiple different load cases and designs must be considered. Kennedy and Martins [22], for example, presented a panel method coupled to a Finite Element Method/Model (FEM) wingbox model, compared various solution methods and investigated the scaling on multiple processors.

Elham and van Tooren [23] presented a low-fidelity optimization framework consisting of a beam FEM model coupled to a Vortex Lattice Method (VLM) aerodynamic model. Six planform and 160 airfoil shape design variables were used to parametrize the geometry. The structural design variables were limited to sizing and the gradients were computed using constraint aggregation in conjunction with the adjoint method.

1.2 SENSITIVITY ANALYSIS

The aforementioned literature has shown what appears to be a disposition to use gradient-based strategies in high-fidelity optimization frameworks. Indeed, the bulk of the work in the previous section is based on gradient-based methods. The main reason can be directly associated to the high computational cost incurred by high-fidelity models. Zeroth-order methods that do not rely on gradients (e.g. genetic algorithms, neural networks, simulated annealing and so forth) provide a better option for locating a global optimum, but they generally require significantly more function evaluations compared to gradient-based methods. Consequently, to minimize the overall computational cost by virtue of reducing the amount of function evaluations, gradient-based methods are often the preferred choice when high-fidelity models are employed. As the name suggests, gradient-based optimization additionally requires sensitivity analysis to determine the derivatives of various functions of interest (e.g. lift, drag, displacements, stresses etc.) with respect to the design variables (e.g. shape, planform, element thicknesses and stiffnesses etc.). Sensitivity analysis is therefore an integral part of gradient-based optimization which enables the optimizer to determine a direction of improvement in the ultimate task of locating an optimum.

Maute et al. [13] computed the gradients of a high-fidelity aero-structural optimization framework by a Linear Block Gauss-Seidel (LBGS) method. This approach is similar to a fixed-point iterative scheme and offers a high level of software modularity. The gradient solvers of the aerodynamic and structural disciplines, if available, can be reused for the computation of the aero-structural gradients. It was later shown in a follow-up paper by Barcelos et al. [24] that the LBGS can lack in convergence robustness and efficiency for strongly coupled aero-structural problems. They presented a novel method, described as a Schur-Krylov (SK) solver, that was shown to improve on the previous shortcomings of the LBGS.

Martins et al. [15] used a lagged-coupled adjoint (LCA) method to obtain the aero-structural derivatives. This method is similar to the LBGS. Only here the discipline sensitivities are computed in a parallel fashion, in contrast to the LBGS which is a sequential method. In fact, the LCA is nothing short of a block Jacobi method. The advantage of the adjoint formulation of the aero-structural gradient problem is that the computational cost becomes proportional to the number of functions and virtually independent of the number of design variables. This is in direct contrast to the direct method, which was applied by Maute et al. [13], where the computational cost is proportional to the number of design variables and nearly independent of the number of functions.

Kenway et al. [17] computed the gradients via the adjoint approach in a closely-coupled fashion by a Generalized Minimal RESidual (GMRES) method. A block diagonal preconditioner was applied to enhance robustness and accelerate conver-

gence. The structural preconditioning block was represented by an exact factorization of the structural stiffness matrix, whereas the aerodynamic preconditioning block was represented by an Incomplete LU-factorization (ILU) of the flux Jacobian. They show that this method converges faster compared to the LBGS. However, the method requires more effort to implement compared to the LBGS for instance. Moreover, the gradients of an aerodynamic grid perturbation must be evaluated for each GMRES iteration as opposed to the LBGS which only requires an aerodynamic grid perturbation during the structural coupling.

The adjoint method is a powerful utility when the number of response functions are few. This makes it ideal in aerodynamic shape optimization where the response functions are typically limited to lift, drag and moment coefficients. Structural optimization, on the other hand, requires that the yield stress on each finite element of the structural model is not exceeded. The direct application of the adjoint method is therefore not suitable for structural models that have hundreds or thousands of elements. To circumvent this predicament, several authors have implemented constraint aggregation techniques that lump all the element stresses into one single constraint function [25, 26]. The benefit of the constraint aggregation is that the adjoint method becomes applicable and can be used to compute the sensitivity derivatives with a substantial reduction in computational time. However, it is widely acknowledged that the constraint aggregation strategy results in suboptimal designs. The structural response is smoothed by the constraint aggregation function and a high level of experience is required to produce a satisfactory outcome by this approach.

An alternative approach of computing high-fidelity derivatives of structural responses in aero-structural design was proposed by Giunta [27]. Modal analysis was employed to reduce the coupling bandwidth between the aerodynamic and structural disciplines. The approximation of the structural deflections by linear superposition of mode shapes makes it affordable for the derivatives to be computed using forward step finite differencing. The aerodynamics in this work was modeled by the Euler equations. However, the choice of the step size is always a concern with standard finite difference schemes and, if not chosen properly, might lead to inaccurate gradients. This very problem of finding an optimal step size was later successfully circumvented by using a complex-step finite difference scheme, as proposed by Martins [28].

A strategy initially developed by Blondeau et al. [29] employed a linearized aerodynamic formulation based on an uncoupled non-intrusive approach. The main advantage of the method is the independency of the computational cost on the number of constraints and potentially on the number of structural design parameters. Moreover, the approach only requires computations of the linearized aerodynamic problem. The method was later demonstrated by Achard et al. [30] on a Common Research Model (CRM) configuration modeled by the RANS equations coupled to a detailed wingbox finite element model.

1.3 RESEARCH OBJECTIVE

Developing a high-fidelity aero-structural optimization framework with limited resources is a massive undertaking and would not add much value or novelty to the current state-of-the-art. A significant effort would have to be allocated to software programming and the required time span would stretch beyond the confines of a doctoral research project. Consequently, the scope of this thesis is limited to target an important facet in the field of high-fidelity aero-structural optimization, namely *sensitivity analysis*. In particular, the focus is narrowed down to structural response gradients in high-fidelity aero-structural optimization, i.e. gradients associated to structural displacements, stresses, buckling and so forth.

To this day, the state-of-the-art of computing structural response gradients in high-fidelity aero-structural optimization frameworks is via the Kreisselmeier-Steinhauser (KS) constraint aggregation function [31]:

$$KS[\mathbf{g}(\mathbf{p})] = g_{max}(\mathbf{p}) + \frac{1}{\rho} \ln \left[\sum_{j=1}^m e^{\rho(g_j(\mathbf{p}) - g_{max}(\mathbf{p}))} \right] \quad (1.1)$$

where m represents the number of constraint functions g , \mathbf{p} is the vector of structural design variables, g_{max} is the maximum constraint evaluated at \mathbf{p} and ρ is the KS parameter. This function offers a great amount of flexibility as it can be applied in conjunction with the adjoint method to compute the gradients at a very low computational cost. Moreover, the KS function is conservative (i.e. it represents a lower bound of all the functions used in the lumping process) and can describe the maximum function in a smooth and differentiable manner. However, the method has two main shortcomings. Firstly, the method provides a global approximation of the response functions. Local and sharp changes in the response derivatives can become difficult to be accurately captured with this method. To account for strong, local variations in the response, which for instance can occur at the intersection of various components such as ribs and wing skins, multiple aggregation functions are often applied on a component basis or on aggregation domains [32]. However, as the number of aggregation domains increases, so does the KS functions and by extension the computation time.

Secondly, and perhaps more importantly, the KS function includes a KS parameter that must be adjusted. Large values (e.g. $\rho \geq 50$) will reduce the distance of the KS envelope to the maximum response value. This typically results in an ill-conditioned optimization problem due to strong gradient variations which can lead to difficult convergence of the optimization. Too small values, on the other hand, condense the design space and result in overly conservative and suboptimal designs. A rudimentary example of the KS function is illustrated in Figure [1.3].

It shows how two inequality constraints f_1 and f_2 change with respect to one design variable x and how the KS function changes with varying values on the KS parameter. The example clearly emphasizes the importance of the KS parameter.

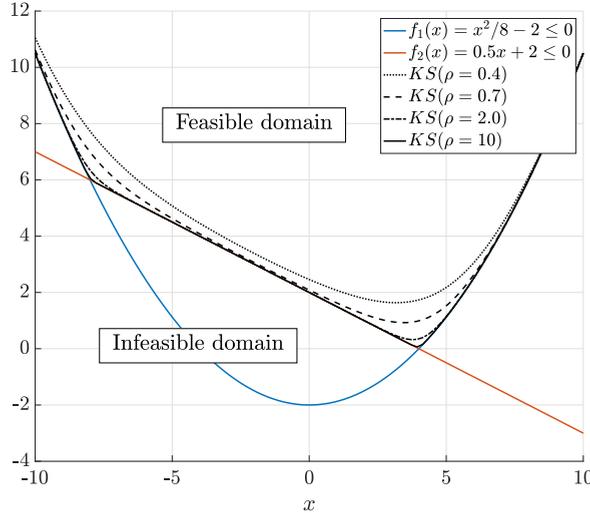


Figure 1.3: Variation of the KS parameter

To determine the optimal value of the KS parameter, there is typically a need for a separate sensitivity study. Alternatively, a value can be assigned based on experience from similar configurations in other optimization problems. $\rho = 50$ is often mentioned as a reasonable value [25, 33].

The limitations described above in the current state-of-the-art has prompted in the formulation of the following two goals for this doctoral research project:

1. *Develop a method and establish a framework for solving direct high-fidelity aero-structural gradient problems.*
2. *Develop a method for computing structural response gradients in high-fidelity aero-structural sensitivity analysis.*

The research in this thesis can be considered an extension of the work by Johannes Dillinger [34], who dedicated his research to develop a sophisticated NASTRAN-based static aeroelastic tailoring framework. The aerodynamic loads computed by the Doublet Lattice Method (DLM) in NASTRAN were corrected by higher fidelity Euler CFD. In the current work, the idea was to not only correct aerodynamic loads, but also investigate the possibility of correcting the corresponding gradients used by the optimizer.

A condition that was imposed on the current research was to incorporate a low/medium-fidelity aerodynamic model in the methodology. Medium-fidelity aerodynamic methods, such as the VLM, have been developed and perfected for many decades and are frequently employed in conceptual design. The idea was to leverage any useful information they might carry in order to accomplish our defined research objectives.

1.4 THESIS LAYOUT

The thesis is divided into five main chapters. The first chapter is the introduction and outlines previous work and current state-of-the-art in MDO and sensitivity analysis. In the second chapter, the first research objective is addressed. A framework is developed for computing structural response gradients in aero-structural sensitivity analysis. The design variables are restricted to structural sizing, i.e. element thicknesses. The framework is divided into modules that interact with one another. It consists of: (i) a low/medium-fidelity aerodynamic solver, (ii) a high-fidelity aerodynamic solver, (iii) a structural solver and (iv) a coupling module that facilitates the load transfer/mesh deformation between the structural and aerodynamic models. A method is presented that extracts information from the low/medium-fidelity aerodynamic model, in particular an aerodynamic stiffness contribution, in order to accelerate the high-fidelity aero-structural gradient computation. The method is demonstrated in a case study with the ONERA M6 (OM6) wing model and its efficiency is compared to a conventional fixed-point iterative scheme.

In the third chapter, the second research objective is addressed. A gradient reconstruction method is presented that approximates high-fidelity aero-structural gradients. The main goal is to be able to compute the gradients of multiple response functions w.r.t. multiple design variables without the convenience of the constraint aggregation approach. A vortex lattice aerodynamic model is used in conjunction with Singular Value Decomposition (SVD) and a pseudo-inverting process in order to reconstruct the high-fidelity aero-structural gradients. The method is demonstrated on the OM6 wing model and a slender Forward-Swept Wing (FSW) configuration.

In chapter four, a static aeroelastic tailoring study is performed on the same FSW configuration to investigate the potential of the gradient reconstruction method described in chapter three. The structural model is parametrized by an in-house conceptual design tool. The model includes ribs, spars, stringers and surface patches but only the spars and the surface patches are included as design variables, whereas ribs and stringers exhibit constant properties throughout the optimization. Two aerodynamic models are applied: (i) a VLM method solving the linearized potential flow equations and (ii) an Euler flow solver. Three

optimization cases are designed to study the convergence behavior when VLM loads and gradients are corrected by the Euler flow model. The first optimization case is a standard structural weight minimization where aerodynamic loads and gradients are computed by the VLM solver. This case is comparable to what was already carried out in the doctoral thesis by Werter [35]. The second case applies Euler load corrections and is comparable to the work by Dillinger [34]. The third case, which is considered a novelty in this thesis, implements not only Euler load corrections, but also the proposed method in chapter three to reconstruct the Euler aero-structural gradients.

The thesis is concluded in the fifth chapter, where all the previous chapters are reviewed and the successfulness of the objectives described in the introduction chapter is assessed. Moreover, suggestions on future work are provided.

2

GRADIENT COMPUTATION

In this chapter, a method is presented that facilitates the computation of high-fidelity aero-structural gradients. The main idea is to improve the spectral properties of complex aero-structural linear systems by constructing a preconditioner based on vortex lattice aerodynamics. The construction of the preconditioner is computationally inexpensive, making it easy to store and invert. The proposed method is tested on an aero-structural gradient computation of the ONERA M6 (OM6) wing. Two high-fidelity aerodynamic models are examined: (i) an inviscid Euler model and (ii) a viscous RANS model with a Spalart-Allmaras turbulence model. The Euler and the RANS grids have 1.05 and 3.57 million cells, respectively. The aerodynamic models are coupled to a linear structural finite element model with 1812 Degrees of Freedom (DOF).

The remainder of this chapter is organized as follows: the computational modules necessary for the gradient computations are discussed in Section [2.1](#). The static aeroelastic analysis, which is a prerequisite for the gradient computation is discussed in Section [2.2](#). The aero-structural gradient computation with the proposed method is discussed in Section [2.3](#), followed by a case study in Section [2.4](#). The chapter is concluded by a brief reflection of the results in Section [2.5](#).

2.1 COMPUTATIONAL MODULES

It is important to preserve a high level of software modularity in work associated with multi-disciplinary analysis and optimization. Organizing the work in modules might increase the overall computation time compared to monolithic codes. However, a higher level of software modularity is coveted in industry as it enables engineers to choose among off-the-shelf solvers with minimal effort required for code adaptation. To this end, the work in this thesis is organized by computational modules, with each module performing individual subtasks and being capable of exchanging information in a systematic fashion. This section provides an overview of the modules and resources necessary for the aeroelastic analysis and the subsequent sensitivity analysis.

In the four subsections that follow the following is outlined: in [2.1.1](#) the aerodynamic solvers, in [2.1.2](#) the structural solver, in [2.1.3](#) the coupling module and in [2.1.4](#) the VLM aerodynamic stiffness matrix.

2.1.1 AERODYNAMIC SOLVERS

Two aerodynamic solvers are necessary in this thesis: (i) one solving a set of low- or medium-fidelity equations and (ii) one solving a set of high-fidelity equations. It is common to associate panel methods that solve the linear potential flow equations with the medium-fidelity category. Singular elements, such as vortex rings or doublets, are used to approximate the flow potential. In this thesis, an in-house VLM code is used as the dedicated medium-fidelity solver [\[36\]](#). The two-dimensional panels are distributed on the camber surface of the wing. The discrete residual of the linear potential flow equations $\tilde{\mathcal{R}}_a$ can be expressed as

$$\tilde{\mathcal{R}}_a(\mathbf{\Gamma}, \mathbf{X}_p) = \mathbf{A}_{IC}\mathbf{\Gamma} + \mathbf{V}_b = \mathbf{0} \quad (2.1)$$

where $\mathbf{\Gamma}$ is the unknown vector of vortex strengths, \mathbf{A}_{IC} is the dense aerodynamic influence coefficient matrix and \mathbf{V}_b is the flow tangency boundary condition vector. Both \mathbf{A}_{IC} and \mathbf{V}_b are functions of the panel coordinates \mathbf{X}_p . Terms associated with the VLM will be denoted by a tilde notation in the remainder of this thesis. The Aerodynamic Influence Coefficient (AIC) matrix can be assembled through the application of the Biot-Savart law. Once the linear system in Equation [2.1](#) is solved, the aerodynamic loads $\tilde{\mathcal{Q}}_a$ can be obtained in a post-treatment step by the Kutta-Joukowski theorem

$$\tilde{\mathcal{Q}}_a = \rho_\infty \mathbf{V}_{tot} \times \mathbf{\Gamma} \quad (2.2)$$

where ρ_∞ is the free-stream density and \mathbf{V}_{tot} is the total velocity vector. Moreover, a Prandtl-Glauert correction is applied to account for compressibility effects in

high-subsonic flows. The theory and the computational routines for the VLM solver are in accordance with, and can be further studied in, Katz and Plotkin [37]. The main purpose of the VLM solver in this chapter is to generate an aerodynamic stiffness matrix, $\tilde{\mathbf{K}}_a$. The derivation of this matrix is described in Subsection 2.1.4 and its purpose will be discussed in Section 2.3.

The main shortcoming of the VLM is its limited application to attached flow conditions in the low-subsonic to moderately high-subsonic regime. The linear potential flow methods lack the predictive capability of modeling nonlinearities in transonic flows. To address this shortcoming, a designated high-fidelity flow solver is required. The high-fidelity solver throughout this thesis is the multi-block structured flow solver ELSA [38]. It is a state-of-the-art industrial CFD code co-developed by ONERA, Airbus and Safran, and is used in the current work to solve the steady-state Euler/RANS equations. In this chapter, the focus is narrowed down to two high-fidelity models of the OM6 wing: (i) one which is governed by the inviscid Euler equations and (ii) one which is governed by the viscous RANS equations depicted in Figures 2.1 and 2.2, respectively. The discrete steady-state flow problem for both types of equations can be expressed by the aerodynamic residual \mathcal{R}_a as

$$\mathcal{R}_a(\mathbf{W}, \mathbf{X}_a) = \mathbf{0} \quad (2.3)$$

where \mathbf{W} is the unknown vector of conservative variables and \mathbf{X}_a are the mesh coordinates. The conservative variables are $(\rho, \rho u, \rho v, \rho w, \rho E)$, where ρ is the density, u , v and w are the Cartesian velocity components and E is the total energy. The unknown vector results from a concatenation of the continuity, the momentum and the energy equations. Hence, the Euler equations have five unknown entries per cell with one or two additional entries for the RANS equations depending on what turbulence model is considered. An implicit Backward-Euler (BWE) time integration scheme is used to advance the flow to a steady-state. The upwind Roe scheme with a Monotonic Upwind Scheme for Conservation Laws (MUSCL) interpolation associated to a Van Albada limiter is applied for the spatial discretization (see Hirsch [39] for instance). Moreover, a one-equation Spalart-Allmaras turbulence model is selected for the RANS problem.

2.1.2 STRUCTURAL SOLVER

The structural solver is the finite element analysis program NASTRAN. Its main task is to generate the global structural stiffness matrix \mathbf{K} to be used externally for the aeroelastic analysis and the sensitivity analysis. If the static deformations are small, the structure can be considered to exhibit linear-elastic behavior. Hence, the resulting stiffness matrix needs only to be generated once for the initial finite element model and stored in memory. In the event of large structural deformations

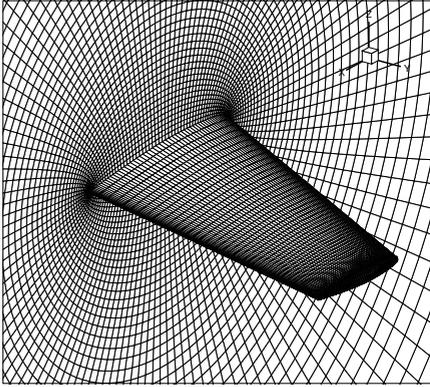


Figure 2.1: OM6 Euler surface grid

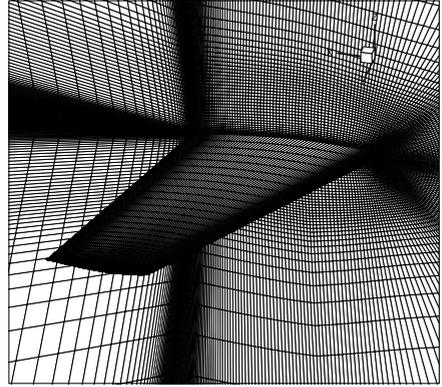


Figure 2.2: OM6 RANS surface grid

the assumption of linear elasticity no longer holds and the stiffness matrix must be updated based on the deformations. In this thesis, only linear elasticity is considered. The static equations can be expressed by the structural residual \mathcal{R}_s :

$$\mathcal{R}_s(\mathbf{U}, \mathbf{X}_s) = \mathbf{K}\mathbf{U} - \mathbf{Q} = \mathbf{0} \quad (2.4)$$

where \mathbf{U} is the vector of structural DOF, \mathbf{X}_s are the structural grid coordinates and \mathbf{Q} is the applied load vector. In addition to the stiffness matrix, pseudo-loads are required for the sensitivity analysis. The pseudo-loads can be derived by differentiating Equation 2.4 with respect to each design variable

$$\frac{d\mathcal{R}_s}{dp_i} = \mathbf{K} \frac{d\mathbf{U}}{dp_i} + \underbrace{\frac{\partial \mathbf{K}}{\partial p_i} \mathbf{U}}_{\text{pseudo-loads}} - \frac{d\mathbf{Q}}{dp_i} = \mathbf{0} \quad (2.5)$$

where the subscript i designates the i^{th} design variable. This equation is derived, and the computation of the terms is explained, in the "Nastran Design Sensitivity and Optimization User's Guide" [40]. What is of interest prior to the sensitivity analysis, and considered a pre-processing step, are gradient terms with a direct dependence to structural sizing variables. The only term in Equation 2.5 that satisfies this is the partial derivative of the stiffness matrix with respect to the design variables

$$\frac{\partial \mathcal{R}_s}{\partial p_i} = \frac{\partial \mathbf{K}}{\partial p_i} \mathbf{U} \quad (2.6)$$

which is obtained by a finite difference scheme. The two remaining terms in Equation 2.5 are implicit functions of the residual, i.e. they can only be obtained by solving the partial differential equations of the linearized aero-structural problem. The term in Equation 2.6, as will be discovered later in Section 2.3, constitutes

the right-hand side in the aero-structural gradient computation and is of size $n_{fdof} \times n_p$, where n_{fdof} is the number of unconstrained structural DOF and n_p is the number of structural design variables. Consequently, the number of right-hand sides is equal to the number of design variables.

The stiffness matrix and the pseudo-loads can be accessed by adding Direct Matrix Abstraction Program (DMAP) [41] alters in the case control section of the NASTRAN bulk data files. To extract the stiffness matrix a SOL 101 static analysis solution sequence is required, whereas the pseudo-load requires a SOL 200 sensitivity analysis solution sequence. The DMAP alters force the extraction of the matrices during the execution of NASTRAN. The exported files in `op4` format are converted to convenient sparse matrix triplet format by external Fortran scripts.

2.1.3 COUPLING MODULE

The purpose of the coupling module is twofold: (i) to determine the aerodynamic grid motion subject to given structural displacements and (ii) to determine the aerodynamic loads acting on the structural mesh points. As the aerodynamic and the structural grids are nonconforming throughout this thesis, there exists a necessity for a robust and efficient mesh deformation algorithm. To this end, an algebraic interpolation routine based on Radial Basis Function (RBF) is applied. The basic principle is to formulate a coupling matrix $[\mathbf{H}]$ that translates structural deformations into aerodynamic grid deformations:

$$\delta \mathbf{X}_a = [\mathbf{H}] \delta \mathbf{X}_s \quad (2.7)$$

where $\delta \mathbf{X}_a$ and $\delta \mathbf{X}_s$ are aerodynamic and structural grid perturbations, respectively. The coupling matrix in this work is derived in accordance with the methodology in Beckert and Wendland [42]. A globally supported Thin-Plate Spline (TPS) function is employed for the Euler and RANS grids which results in a dense interpolation matrix. To reduce the memory usage, the coupling matrix is only formulated as a matrix-vector product. Moreover, the mesh deformation is executed in parallel mode on multiple processors. As the RBF approach does not require any connectivity information, parallelizing the code becomes a trivial task.

A simple test case is conducted to examine the parallel computing scalability. Mesh deformation is performed on the 1.05 million cell OM6 Euler model. The computations are performed on 2.60 GHz Intel Xeon E5-2660 v3 Processors. The wall time is clocked for 2, 4, 8 and 16 processors. The results in Figure 2.3 indicate very good scalability up to 16 processors. The wall time for one mesh deformation on a 1.05 million cell model can be reduced from 56 seconds on 2 processors to 9 seconds on 16 processors, which is more than a sixfold improvement.

Another method of reducing the computation time, in addition to code parallelism, is to limit the number of surface points in the formulation of the coupling

matrix. There are numerous data point reduction algorithms that can be used to condense the size of the coupling matrix without significantly affecting the quality of the deformation [43, 44]. In this chapter, due to the structured mesh, a subset of every 9th surface grid point is used in the formulation of the coupling matrix, treating the remaining grid points as unknowns. This subset has been found to yield a favorable trade-off between mesh deformation accuracy and computation time. The selected surface points for the Euler grid are depicted in Figure 2.4.

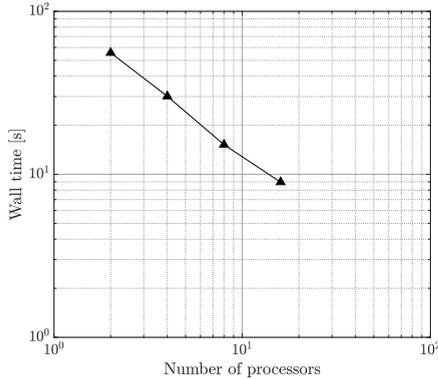


Figure 2.3: Parallel computing scalability

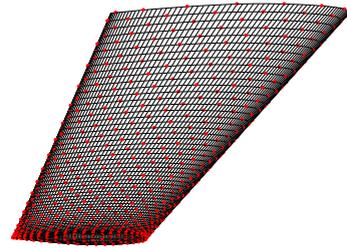


Figure 2.4: Selected surface points for mesh deformation

In previous work [45], the equivalence of virtual work [42] was relied upon, which allows for the transposed coupling matrix to be used as a load transfer operator

$$\mathbf{Q} = [\mathbf{H}]^T \mathbf{Q}_a \quad (2.8)$$

where \mathbf{Q}_a are loads acting on the aerodynamic surface grid and \mathbf{Q} are the transferred loads on the structure. However, this is not always a favorable option. Based on our experience, this operator can result in unrealistic oscillatory loads on the structural grid for certain configurations. We therefore resort to a simple, yet highly effective and robust, nearest-neighbor scheme for multivariate interpolation [46]. In this thesis, the load transfer operator is denoted by $[\mathbf{T}]$, such that

$$\mathbf{Q} = [\mathbf{T}] \mathbf{Q}_a \quad (2.9)$$

The load transfer operator is constructed by looping through the aerodynamic surface grid points and assigning the loads to the nearest structural point. The moments $[M^x, M^y, M^z]^T$ on the structural points are evaluated by multiplying

the load with Euclidean distance between the points as:

$$\begin{aligned} M^x &= Q_a^z dy - Q_a^y dz \\ M^y &= Q_a^x dz - Q_a^z dx \\ M^z &= Q_a^y dx - Q_a^x dy \end{aligned} \quad (2.10)$$

2.1.4 VLM AERODYNAMIC STIFFNESS MATRIX

As emphasized in Subsection 2.1.1, the aerodynamic stiffness matrix $\tilde{\mathbf{K}}_a$ needs to be provided by the VLM solver. Its derivation is covered in this subsection, whereas its purpose will be explained in Section 2.3. In essence, the VLM aerodynamic stiffness matrix describes how the aerodynamic loads applied on the structural model change due to perturbations of the structural DOF. It can be derived by formulating the static aeroelastic problem, given by:

$$\begin{pmatrix} \tilde{\mathcal{R}}_a(\mathbf{\Gamma}, \mathbf{U}) \\ \mathcal{R}_s(\mathbf{U}, \mathbf{\Gamma}) \end{pmatrix} = \mathbf{0} \quad (2.11)$$

where \mathbf{X}_p and \mathbf{X}_s have been omitted for the sake of brevity. If the aeroelastic residual is linearized with respect to the unknown variables $\mathbf{\Gamma}$ and \mathbf{U} , the following linear system is obtained:

$$\begin{bmatrix} \frac{\partial \tilde{\mathcal{R}}_a}{\partial \mathbf{\Gamma}} & \frac{\partial \tilde{\mathcal{R}}_a}{\partial \mathbf{U}} \\ \frac{\partial \mathcal{R}_s}{\partial \mathbf{\Gamma}} & \frac{\partial \mathcal{R}_s}{\partial \mathbf{U}} \end{bmatrix} \begin{pmatrix} \Delta \mathbf{\Gamma} \\ \Delta \mathbf{U} \end{pmatrix} = - \begin{pmatrix} \tilde{\mathcal{R}}_a \\ \mathcal{R}_s \end{pmatrix} \quad (2.12)$$

where the increments of the vortex strengths and the structural deformations are the Newton corrections. This in fact amounts to Newton's method for solving the static aeroelastic problem given that the partial derivatives in the coefficient matrix are exact and the linear system is solved exactly. The diagonal blocks of the coefficient matrix are single-disciplinary and are therefore determined individually by their respective solvers. The off-diagonal blocks include coupling matrices that require information from both solvers as well as the coupling module. The linear system given by Equation 2.12 can be expressed in Schur complement form by eliminating the aerodynamic degrees of freedom

$$\left(\frac{\partial \mathcal{R}_s}{\partial \mathbf{U}} - \underbrace{\frac{\partial \mathcal{R}_s}{\partial \mathbf{\Gamma}} \left[\frac{\partial \tilde{\mathcal{R}}_a}{\partial \mathbf{\Gamma}} \right]^{-1} \frac{\partial \tilde{\mathcal{R}}_a}{\partial \mathbf{U}}}_{\tilde{\mathbf{K}}_a} \right) \Delta \mathbf{U} = -\mathcal{R}_s \quad (2.13)$$

where the coefficient matrix is the Schur complement and $\tilde{\mathbf{K}}_a$ is the aerodynamic stiffness matrix. This matrix is typically explicitly constructed in VLM codes.

The three blocks required to construct the aerodynamic stiffness matrix are explained next:

- $\partial\mathcal{R}_s/\partial\mathbf{\Gamma}$ describes how the structural residual changes with respect to the vortex strength perturbations. This matrix can be conveniently decomposed as: $(\partial\mathcal{R}_s/\partial\tilde{\mathbf{Q}}_a)(\partial\tilde{\mathbf{Q}}_a/\partial\mathbf{\Gamma})$. The first term is the load transfer operator $[\mathbf{T}]$ that interpolates loads from the aerodynamic panels on the structural nodes. The second term can be easily obtained by differentiating the Kutta-Joukowski theorem, given by Equation 2.2.
- $\partial\tilde{\mathcal{R}}_a/\partial\mathbf{\Gamma}$ describes how the aerodynamic residuals change with respect to the vortex strength perturbations. This is the dense \mathbf{A}_{IC} matrix which is already available in Equation 2.1. This term can be assembled in a straightforward manner by the VLM solver, inverted and stored in memory.
- $\partial\tilde{\mathcal{R}}_a/\partial\mathbf{U}$ describes the linearization of the aerodynamic residuals with respect to the structural DOF. It can be expressed as: $(\partial\tilde{\mathcal{R}}_a/\partial\mathbf{X}_p)(d\mathbf{X}_p/d\mathbf{U})$. The first term requires a linearization of the aerodynamic residual 2.1 with respect to the panel coordinates. Both the AIC matrix and the flow tangency boundary condition vector are direct functions of \mathbf{X}_p . This term is somewhat cumbersome to obtain and can result in large computational overheads for aerodynamic models that require fine panel grids. To partially overcome this shortcoming, a parallel approach with analytically obtained gradients is implemented. The second term in the decomposition is a displacement interpolation operator which is readily available from the coupling module. For the VLM code, an Infinite-Plate Spline (IPS) 47 method is applied to adjust the panels with respect to the structural displacements.

2.2 AEROELASTIC ANALYSIS

The state of equilibrium for a wing aero-structural configuration is determined by a static aeroelastic analysis. This part is crucial in design optimization as it allows for important quantities of interest, also known as responses, to be evaluated. Lift, drag and moment coefficients are typically of main interest in aerodynamic optimization. In the context of structural optimization, quantities such as strains/stresses, displacements and weight are generally of interest. Any of these quantities \mathcal{J} , aerodynamic or structural, may be expressed as a function of a design variable p_i in the form

$$\mathcal{J} = \mathcal{J}(p_i, \mathbf{W}(p_i), \mathbf{U}(p_i)) \quad (2.14)$$

The function of interest can either represent an objective or a constraint. The design variable p_i can impose an explicit influence on the function through geo-

metric shape variations, as well as an implicit influence through the solution of the underlying governing equations. To obtain the function value for a given set of design variables, the nonlinear static aeroelastic problem must be solved. It can be formulated by concatenating the flow problem [2.3](#) and the structural problem [2.4](#) such that

$$\begin{pmatrix} \mathcal{R}_a(\mathbf{W}, \mathbf{U}) \\ \mathcal{R}_s(\mathbf{U}, \mathbf{W}) \end{pmatrix} = \mathbf{0} \quad (2.15)$$

The design variables are omitted here for the sake of brevity as they are constant throughout the analysis. In practice, the concatenated residuals are not reduced to exactly zero. Rather, the coupled nonlinear problem is considered to be solved when there is a sufficient drop in the Euclidean norm of the respective residuals between two consecutive iterations. There are numerous methods for solving the nonlinear problem described by Equation [2.15](#) [\[13, 24, 22, 17\]](#). Throughout this thesis, a defect-correction approach is applied with the basic principles outlined by a flow chart in [Figure 2.5](#).

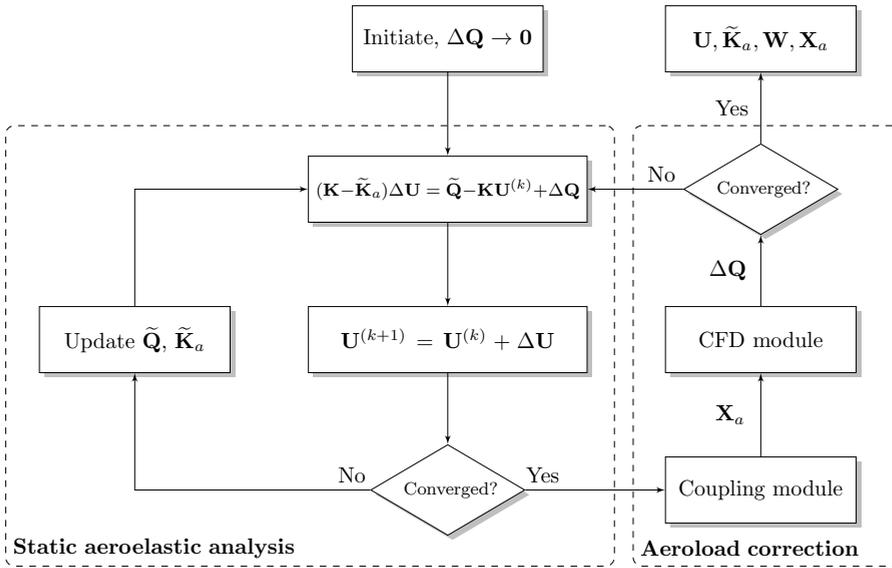


Figure 2.5: The defect-correction approach

A key feature of the method is to apply a defect aeroload, which is defined as the difference between the aerodynamic loads from an Euler/RANS model and those from a VLM model for a given deflection: $\Delta \mathbf{Q} = \mathbf{Q} - \tilde{\mathbf{Q}}$. The method is initiated by setting the defect aeroload vector to zero and then solving the static aeroelastic problem by Newton's method. After the solution has converged (typically after 3-4 iterations), an external aeroload correction module is activated. First the CFD

mesh is deformed based on the structural deflection, before the nonlinear steady-state flow problem is approximately solved and the defect aeroload is estimated. A residual drop of 1-2 order of magnitude is typically sufficient here. It is important to mention that for the subsequent aeroload corrections, the flow is restarted from the previously converged flow solution. As the defect aeroload is continuously updated, it becomes clear that the VLM aeroload in Newton's equation will be canceled out by the VLM term in the defect aeroload. Hence, upon convergence, the static aeroelastic solution will be solely determined by the CFD loads. The feasibility of the static aeroelastic solution is estimated by two criteria and is monitored in the aeroload correction module: (i) a sufficient overall drop of the flow residual and (ii) a sufficient drop in the norm of the difference between two consecutive aeroload corrections. In this chapter, a minimum drop of six orders of magnitude is required for both criteria.

2.3 SENSITIVITY ANALYSIS

Sensitivity analysis plays an important role in gradient-based optimization. The gradients obtained by the sensitivity analysis enable the optimizer to determine a direction of improvement in the ultimate task of finding an optimum. The main goal here is to evaluate the variation of one or several functions of interest due to a perturbation of the i^{th} design variable $d\mathcal{J}/dp_i$. There are several methods for obtaining the total gradient as is discussed by Peter and Dwight [48]. The simplest method to implement is finite difference approximation. For instance, the total gradient can be obtained by the forward finite difference formula:

$$\frac{d\mathcal{J}}{dp_i} \approx \frac{\mathcal{J}(p_i + \Delta p_i, \mathbf{W}(p_i + \Delta p_i), \mathbf{U}(p_i + \Delta p_i)) - \mathcal{J}(p_i, \mathbf{W}(p_i), \mathbf{U}(p_i))}{\Delta p_i} \quad (2.16)$$

where Δp_i is the finite design variable perturbation. A major advantage of this method is that it does not require any modification to the primary solvers. However, it is computationally expensive for large scale problems with many design variables. Forward differencing requires $(n_p + 1)$ function evaluations. Moreover, the gradients can be very sensitive to the magnitude of the perturbation. Too large perturbations result in inaccurate gradients, whereas too small perturbations result in subtractive cancellation. Another method of obtaining the total gradient, which will be the focus in this work, is based on analytical differentiation of the function of interest. Applying the chain rule to Equation 2.14 yields

$$\frac{d\mathcal{J}}{dp_i} = \frac{\partial \mathcal{J}}{\partial \mathbf{W}} \frac{d\mathbf{W}}{dp_i} + \frac{\partial \mathcal{J}}{\partial \mathbf{U}} \frac{d\mathbf{U}}{dp_i} \quad (2.17)$$

In this chapter, the scope is reduced to structural sizing variables. Hence, partial derivatives of the structural or aerodynamic grid with respect to the design

variables ($\partial \mathbf{X}_s / \partial p_i$ and $\partial \mathbf{X}_a / \partial p_i$) are equal to zero, e.g. a change in the thickness of a structural member does not have an explicit impact on the wing shape. The partial derivatives of the function of interest with respect to the state variables, $\partial \mathcal{J} / \partial \mathbf{W}$ and $\partial \mathcal{J} / \partial \mathbf{U}$, are generally straightforward to evaluate, at least for quantities related to the wall pressure field (lift coefficient or pressure drag coefficient), and can be provided externally by a pre-processing step. However, to obtain the total gradients of the state variables, $d\mathbf{W} / dp_i$ and $d\mathbf{U} / dp_i$ the discrete static aeroelastic problem [2.15](#) must be differentiated with respect to the design variable p_i :

$$\begin{pmatrix} \frac{d\mathcal{R}_a}{dp_i} \\ \frac{d\mathcal{R}_s}{dp_i} \end{pmatrix} = \begin{pmatrix} \frac{\partial \mathcal{R}_a}{\partial p_i} \\ \frac{\partial \mathcal{R}_s}{\partial p_i} \end{pmatrix} + \begin{bmatrix} \frac{\partial \mathcal{R}_a}{\partial \mathbf{W}} & \frac{\partial \mathcal{R}_a}{\partial \mathbf{U}} \\ \frac{\partial \mathcal{R}_s}{\partial \mathbf{W}} & \frac{\partial \mathcal{R}_s}{\partial \mathbf{U}} \end{bmatrix} \begin{pmatrix} \frac{d\mathbf{W}}{dp_i} \\ \frac{d\mathbf{U}}{dp_i} \end{pmatrix} = \mathbf{0} \quad (2.18)$$

Note that the total gradient of the residuals is zero. This is consistent with the assumption of a zero residual in Equation [2.15](#), and is a prerequisite for accurate gradients. If the solution vector $(d\mathbf{W} / dp_i \ d\mathbf{U} / dp_i)^T$ is solved for and subsequently substituted in Equation [2.17](#), the total gradient of the function of interest can be expressed conceptually as

$$\frac{d\mathcal{J}}{dp_i} = - \begin{pmatrix} \frac{\partial \mathcal{J}}{\partial \mathbf{W}} & \frac{\partial \mathcal{J}}{\partial \mathbf{U}} \end{pmatrix} \begin{bmatrix} \frac{\partial \mathcal{R}_a}{\partial \mathbf{W}} & \frac{\partial \mathcal{R}_a}{\partial \mathbf{U}} \\ \frac{\partial \mathcal{R}_s}{\partial \mathbf{W}} & \frac{\partial \mathcal{R}_s}{\partial \mathbf{U}} \end{bmatrix}^{-1} \begin{pmatrix} \frac{\partial \mathcal{R}_a}{\partial p_i} \\ \frac{\partial \mathcal{R}_s}{\partial p_i} \end{pmatrix} \quad (2.19)$$

There are in principle two approaches for computing the total gradient. The first, and probably more intuitive approach, is to solve the linear system as presented in Equation [2.18](#) and then multiply the total derivatives of the conservative variables and structural deformations by the functional response partial derivatives. This approach is known as the *direct approach*. Since the bulk of the computation time is consumed by the solution of the linear system, it is easy to anticipate a computation time that scales linearly with the number of right-hand sides which in this instance is the number of design variables. The alternative approach is to solve the left linear system in Equation [2.19](#) to obtain the aero-structural adjoint vector

$$\begin{bmatrix} \frac{\partial \mathcal{R}_a}{\partial \mathbf{W}} & \frac{\partial \mathcal{R}_a}{\partial \mathbf{U}} \\ \frac{\partial \mathcal{R}_s}{\partial \mathbf{W}} & \frac{\partial \mathcal{R}_s}{\partial \mathbf{U}} \end{bmatrix}^T \begin{pmatrix} \lambda_a \\ \lambda_s \end{pmatrix} = \begin{pmatrix} \frac{\partial \mathcal{J}}{\partial \mathbf{W}} & \frac{\partial \mathcal{J}}{\partial \mathbf{U}} \end{pmatrix}^T \quad (2.20)$$

where λ_a and λ_s are the aerodynamic and structural adjoint vectors, respectively. The total gradient can then be obtained in a straightforward manner by the matrix-vector multiplication

$$\frac{d\mathcal{J}}{dp_i} = - (\lambda_a \quad \lambda_s) \begin{pmatrix} \frac{\partial \mathcal{R}_a}{\partial p_i} \\ \frac{\partial \mathcal{R}_s}{\partial p_i} \end{pmatrix} \quad (2.21)$$

This approach is known as the *adjoint approach*. Contrary to the direct approach, the computation time of the adjoint approach scales linearly with the number of functions. It can be concluded that the adjoint approach is a more efficient method when the number of design variables exceed the number of functions, and conversely the direct approach is a more efficient method when the number of functions exceed the number of design variables.

Before any solution methods of the linear systems are discussed, it is important to understand the partial derivatives in the coefficient matrix in Equation 2.19 and what implications they might bring about. They are outlined in the following:

2

- $\partial \mathcal{R}_a / \partial \mathbf{W}$ describes how the flux residuals change with respect to flow variables. It is commonly referred to as a flux Jacobian and can be obtained by linearizing the discrete residual in Equation 2.3 with respect to the conservative variables. For the Euler equations, this term represents an exact linearization of the explicit residual. However, because of the complexity in achieving an accurate linearization for the RANS equations, and also in order to significantly reduce the required level of effort, it is common to consider simplifying assumptions such as the "frozen" turbulence model approach. In addition, the exact linearization of the turbulence model leads to high condition number of the flux Jacobian matrix which essentially makes the linear system difficult to solve.

When the "frozen- μ_t " assumption is applied, only the first five residual equations are linearized with respect to the mean-flow variables. The eddy viscosity μ_t is simply treated as a constant. It should be mentioned that the accuracy of the obtained gradients are directly dependent on the accuracy of the linearization of this term. Nevertheless, a limited number of authors have considered turbulence model linearization and discussed its benefit in terms of accuracy [49, 50, 51].

- $\partial \hat{\mathcal{R}}_a / \partial \mathbf{W}$ is not explicitly stated in the coefficient matrix. However, due to its importance, it is worth mentioning. The exact Jacobian matrix (or its transpose) appearing in Equations 2.19 and 2.20 has to be inverted in some way. To avoid this complex task, those equations are put in an incremental iterative form and typically solved by an implicit method (BWE scheme). $\partial \hat{\mathcal{R}}_a / \partial \mathbf{W}$ represents an accurate approximation of the flux Jacobian that appears on the left-hand side (the implicit stage) and the solution increment is obtained by approximately solving the linear system using few steps of a Lower-Upper Symmetric Successive OverRelaxation (LU-SSOR) scheme [52].

The approximate Jacobian is usually a standard first-order upwind approximation of the differential of the flux balance which is preferred for better numerical conditioning. In contrast to the flux Jacobian, the accuracy of

the obtained gradients are not dependent on the accuracy of this term. This term only affects the convergence rate of the gradient computation.

- $\partial\mathcal{R}_a/\partial\mathbf{U}$ describes how the flow residual changes due to perturbations on the structural DOF. This term can be expressed as a product of two differentials: $(\partial\mathcal{R}_a/\partial\mathbf{X}_a)(d\mathbf{X}_a/d\mathbf{U})$. The second term is the mesh deformation operator $[\mathbf{H}]$ which was covered in Subsection 2.1.3. The first term is a geometric sensitivity of the explicit residual. In contrast to the flux Jacobian, this term is linearized by fixing the flow variables, while perturbing the aerodynamic grid. Due to the limitations of the ELSA version that is used in this thesis, this particular term is obtained as a matrix-vector product by central differencing. Nevertheless, finite difference approximations of this particular term are known to be robust [53]. The matrix-vector product is computed by:

$$\frac{\partial\mathcal{R}_a}{\partial\mathbf{X}_a} \frac{d\mathbf{X}_a}{dp_i} \approx \frac{\mathcal{R}_a\left(\mathbf{W}(p_i), \mathbf{X}_a(p_i) + \Delta p_i \frac{d\mathbf{X}_a}{dp_i}\right) - \mathcal{R}_a\left(\mathbf{W}(p_i), \mathbf{X}_a(p_i) - \Delta p_i \frac{d\mathbf{X}_a}{dp_i}\right)}{2\Delta p_i} \quad (2.22)$$

where $d\mathbf{X}_a/dp_i = [\mathbf{H}](d\mathbf{U}/dp_i)$ is provided as input to the solver. Note that the conservative variables are kept constant when the residual is recalculated for the perturbed grid. The finite perturbation magnitude Δp_i is set to $1e-3$.

- $\partial\mathcal{R}_s/\partial\mathbf{W}$ describes how the structural residual in Equation 2.4 changes due to flow variable perturbations. The only term that is directly influenced by the conservative variables are the aerodynamic loads on the structural grid. The term can be decomposed as the following product of three differentials: $(\partial\mathbf{Q}/\partial\mathbf{Q}_a)(\partial\mathbf{Q}_a/\partial\mathbf{W}_b)(\partial\mathbf{W}_b/\partial\mathbf{W})$ where \mathbf{W}_b are boundary flow variables that satisfy the physical boundary conditions. The first term is the load transfer operator $[\mathbf{T}]$ that translates loads and moments from the aerodynamic boundary to the structural DOF. The second term is obtained by linearizing the surface load pressure integration scheme with respect to the wall flow variables. Only cells adjacent to the surface affect this term. The third term extrapolates the flow variables based on the numerical scheme from the adjacent cell centers to the boundary where the loads are calculated.
- $\partial\mathcal{R}_s/\partial\mathbf{U}$ describes how the structural residual in Equation 2.4 changes due to the structural DOF. This term includes the global structural stiffness matrix \mathbf{K} and the geometric aerodynamic load gradient $\partial\mathbf{Q}/\partial\mathbf{U}$. The first term was already covered in Subsection 2.1.2. The second term can be decomposed as a product of three differentials: $(\partial\mathbf{Q}/\partial\mathbf{Q}_a)(\partial\mathbf{Q}_a/\partial\mathbf{X}_a)(d\mathbf{X}_a/d\mathbf{U})$. The first and the third term are the load transfer operator $[\mathbf{T}]$ and mesh deformation operator $[\mathbf{H}]$, respectively. The second term describes how the

loads on the boundary change due to perturbations on the grid coordinates. It can be obtained by linearizing the surface load pressure integration scheme with respect to the surface normals.

An interesting prospect is to solve the linear problem in Equation 2.19, either by adjoint or direct mode, in a closely coupled fashion using a Krylov-based solver. Kenway et al. [17] solve the aero-structural gradient problem in adjoint mode using GMRES and the solution method was referred to as a Coupled-Krylov solver. They show that this method exhibits faster convergence rates compared to more conventional methods. However, solving this coupled linear system is not a trivial task and presents a number of challenges. Firstly, the scaling between the unknowns from the different disciplines must be accounted for. If this is not the case, the norm of the residual which is used to determine the acquired level of convergence will not be representative for both disciplines and the desired level of accuracy might not be obtained as discussed by Wathen [54]. Secondly, the cross-coupling term $\partial \mathcal{R}_a / \partial \mathbf{U} \equiv (\partial \mathcal{R}_a / \partial \mathbf{X}_a)[\mathbf{H}]$ in Equation 2.19 implies that a mesh deformation computation is required for each GMRES iteration. This results in a total gradient solution time that is highly dependent on the efficiency of the mesh deformation module. For RANS problems, high-quality mesh deformation methods may be required at times to preserve boundary layer orthogonality [55] and ensure properly converged flow solutions. In such cases, the total gradient computation time may even become dominated by the mesh deformation gradient. Finally, it should be mentioned that a closely coupled or monolithic solution procedure, albeit efficient and robust, requires higher implementation costs which can diminish the software modularity.

Alternatively, Equation 2.19 can be solved by separating the coupled problem into two subproblems and solving for each subproblem sequentially. This results in the following formulation for the direct approach:

$$\begin{cases} \frac{\partial \mathcal{R}_a}{\partial \mathbf{W}} \frac{d\mathbf{W}}{dp_i}^{(k+1)} = -\frac{\partial \mathcal{R}_a}{\partial p_i} - \frac{\partial \mathcal{R}_a}{\partial \mathbf{U}} \frac{d\mathbf{U}}{dp_i}^{(k)} & (2.23a) \\ \frac{\partial \mathcal{R}_s}{\partial \mathbf{U}} \frac{d\mathbf{U}}{dp_i}^{(k+1)} = -\frac{\partial \mathcal{R}_s}{\partial p_i} - \frac{\partial \mathcal{R}_s}{\partial \mathbf{W}} \frac{d\mathbf{W}}{dp_i}^{(k+1)} & (2.23b) \end{cases}$$

The flow subproblem in Equation 2.23a is approximately solved before the structural subproblem in Equation 2.23b is solved with a right-hand side that includes the latest update of the aerodynamic unknowns. Hence, this method can be regarded as a LBGS method and follows the same line of reasoning as in Maute et al. [13]. It is a sequential method since only one subproblem can be solved at a time. The total solution is obtained when a significant drop in the residual norm of the two subproblems is observed. The LBGS problem above can be reformulated with the terms described previously in this work as follows:

$$\left\{ \begin{array}{l} \frac{\partial \mathcal{R}_a}{\partial \mathbf{W}} \frac{d\mathbf{W}}{dp_i}^{(k+1)} = -\frac{\partial \mathcal{R}_a}{\partial \mathbf{X}_a} [\mathbf{H}] \frac{d\mathbf{U}}{dp_i}^{(k)} \end{array} \right. \quad (2.24a)$$

$$\left\{ \begin{array}{l} \mathbf{K} \frac{d\mathbf{U}}{dp_i}^{(k+1)} = -\frac{\partial \mathbf{K}}{\partial p_i} \mathbf{U} + [\mathbf{T}] \left(\frac{\partial \mathbf{Q}_a}{\partial \mathbf{W}} \frac{d\mathbf{W}}{dp_i}^{(k+1)} + \frac{\partial \mathbf{Q}_a}{\partial \mathbf{X}_a} [\mathbf{H}] \frac{d\mathbf{U}}{dp_i}^{(k)} \right) \end{array} \right. \quad (2.24b)$$

Note that $\partial \mathcal{R}_a / \partial p_i$ is dropped from the flow subproblem since only structural sizing is considered in this work. An advantage of the LBGS method is that the gradient computation routines for the flow and structural solver, if existent, can be reused. Only the right-hand sides must be modified accordingly. The disadvantage is that the approach lacks robustness and efficiency for strongly coupled FSI problems as highlighted by Barcelos et al. [24]. These types of problems typically occur when the aerodynamic loads are large compared to the counteracting internal stiffness, which in turn result in large structural deformations.

A different approach can be taken by reformulating the linear system in Equation 2.18 using a Schur complement reduction, effectively eliminating the aerodynamic unknowns $d\mathbf{W}/dp_i$ such that

$$[\mathbf{K} - \mathbf{K}_a] \frac{d\mathbf{U}}{dp_i} = -\frac{\partial \mathbf{K}}{\partial p_i} \mathbf{U} \quad (2.25)$$

where the coefficient matrix $[\mathbf{K} - \mathbf{K}_a]$ is the Schur complement of the block $\partial \mathcal{R}_a / \partial \mathbf{W}$ and \mathbf{K}_a is the high-fidelity aerodynamic stiffness matrix. This matrix can conceptually be constructed by the following sequence of differentials

$$\mathbf{K}_a = [\mathbf{T}] \left(-\frac{\partial \mathbf{Q}_a}{\partial \mathbf{W}} \left[\frac{\partial \mathcal{R}_a}{\partial \mathbf{W}} \right]^{-1} \frac{\partial \mathcal{R}_a}{\partial \mathbf{X}_a} + \frac{\partial \mathbf{Q}_a}{\partial \mathbf{X}_a} \right) [\mathbf{H}] \quad (2.26)$$

It becomes obvious that the aerodynamic stiffness matrix can not be explicitly formulated for realistic problems since it requires the exact inverse of the flux Jacobian $\partial \mathcal{R}_a / \partial \mathbf{W}$. This would result in excessive amounts of computational resources and is consequently not a viable option. In practice, an iterative solution method is applied to Equation 2.25 where only a matrix-vector product of the aerodynamic stiffness matrix is required. For the sake of clarity, the linear problem can be expressed as

$$\mathbf{A} \mathbf{x} = \mathbf{b} \quad (2.27)$$

where

$$\mathbf{A} = [\mathbf{K} - \mathbf{K}_a], \quad \mathbf{x} = \frac{d\mathbf{U}}{dp_i}, \quad \mathbf{b} = -\frac{\partial \mathbf{K}}{\partial p_i} \mathbf{U} \quad (2.28)$$

The idea is to apply an iterative scheme by adopting a standard fixed-point formulation

$$\mathbf{M}\mathbf{x}^{(k+1)} = \mathbf{N}\mathbf{x}^{(k)} + \mathbf{b} \quad (2.29)$$

where the system matrix \mathbf{A} is decomposed as $\mathbf{A} = \mathbf{M} - \mathbf{N}$ and the current iteration is k . Here, \mathbf{M} is a preconditioner that is typically an approximation of the coefficient matrix \mathbf{A} . The assumption is that the preconditioner is easy to invert compared to the coefficient matrix. It can be shown that the convergence characteristics of Equation 2.29 depend on the spectral radius of the iteration matrix \mathbf{G} [56]:

$$\mathbf{G} = \mathbf{I} - \mathbf{M}^{-1}\mathbf{A} \quad (2.30)$$

where \mathbf{I} is the identity matrix. This highlights the importance of a quality preconditioner to improve the convergence behavior. In extension, this implies that the robustness and efficiency of an aero-structural gradient solution scheme is highly influenced by the choice of the preconditioner. The typical procedure is to use an exact factorization of the global structural stiffness matrix, such that $\mathbf{M} \equiv \mathbf{K}$. Equation 2.25 can then be expressed in fixed-point form as

$$\mathbf{K} \left(\frac{d\mathbf{U}^{(k+1)}}{dp_i} - \frac{d\mathbf{U}^{(k)}}{dp_i} \right) = -[\mathbf{K} - \mathbf{K}_a] \frac{d\mathbf{U}^{(k)}}{dp_i} - \frac{\partial \mathbf{K}}{\partial p_i} \mathbf{U} \quad (2.31)$$

The solution converges when the right-hand side goes to zero. Note that this formulation is identical to the LBGs method. It is evident that the preconditioner, or the coefficient matrix in Equation 2.31, does not account for the aerodynamic portion of the stiffness. The proposal in this work is to improve the spectral properties of the preconditioned linear system using the vortex lattice model. Our preconditioning matrix can then be expressed as $\mathbf{M} \equiv [\mathbf{K} - \tilde{\mathbf{K}}_a]$, where $\tilde{\mathbf{K}}_a$ is the contribution from the vortex lattice model. Equation 2.31 can then be replaced by

$$[\mathbf{K} - \tilde{\mathbf{K}}_a] \left(\frac{d\mathbf{U}^{(k+1)}}{dp_i} - \frac{d\mathbf{U}^{(k)}}{dp_i} \right) = -[\mathbf{K} - \mathbf{K}_a] \frac{d\mathbf{U}^{(k)}}{dp_i} - \frac{\partial \mathbf{K}}{\partial p_i} \mathbf{U} \quad (2.32)$$

The contribution from the vortex lattice model is expected to improve the stability and convergence rate when the Schur complement on the right-hand side is ill-conditioned. On the other hand, when the structural stiffness matrix has a dominating influence, for instance when the aeroelastic effect is negligible, the addition of $\tilde{\mathbf{K}}_a$ is not expected to have any notable impact. It should be added that, as stated above, the idea is to improve the spectral properties of the linear problem. However, this can not be quantified by computing the eigenvalues of

the Schur complement since this matrix is never explicitly formulated. This assumption can only be evaluated implicitly by comparing convergence results with and without the addition of $\tilde{\mathbf{K}}_a$. The solution procedure of the proposed method is outlined in detail in Table [2.1](#).

Table 2.1: Proposed method

1. Evaluate $\frac{\partial \mathbf{K}}{\partial p_i} \mathbf{U}$ after the static aeroelastic analysis has converged.
2. Obtain an initial guess of the solution vector $\frac{d\mathbf{U}}{dp_i}^{(0)}$ by solving the medium-fidelity problem:

$$[\mathbf{K} - \tilde{\mathbf{K}}_a] \frac{d\mathbf{U}}{dp_i}^{(0)} = -\frac{\partial \mathbf{K}}{\partial p_i} \mathbf{U}$$
3. Start the iteration with $k = 0$.
4. Formulate the right-hand side vector \mathbf{rhs}_a of the flow subproblem in Equation [2.24a](#):

$$\mathbf{rhs}_a^{(k)} = -\frac{\partial \mathcal{R}_a}{\partial \mathbf{X}_a} [\mathbf{H}] \frac{d\mathbf{U}}{dp_i}^{(k)}$$
5. Approximately solve the flow subproblem:

$$\frac{\partial \mathcal{R}_a}{\partial \mathbf{W}} \frac{d\mathbf{W}}{dp_i}^{(k+1)} = \mathbf{rhs}_a^{(k)}$$

This linear system is solved by GMRES and the process is terminated when the residual norm drops one order of magnitude.
6. Formulate the right-hand side vector \mathbf{rhs}_s of the structural subproblem in Equation [2.32](#):

$$\mathbf{rhs}_s^{(k)} = -\mathbf{K} \frac{d\mathbf{U}}{dp_i}^{(k)} + [\mathbf{T}] \left(\frac{\partial \mathcal{Q}_a}{\partial \mathbf{W}} \frac{d\mathbf{W}}{dp_i}^{(k+1)} + \frac{\partial \mathcal{Q}_a}{\partial \mathbf{X}_a} [\mathbf{H}] \frac{d\mathbf{U}}{dp_i}^{(k)} \right) - \frac{\partial \mathbf{K}}{\partial p_i} \mathbf{U}$$
7. Exactly solve the structural subproblem:

$$[\mathbf{K} - \tilde{\mathbf{K}}_a] \left(\frac{d\mathbf{U}}{dp_i}^{(k+1)} - \frac{d\mathbf{U}}{dp_i}^{(k)} \right) = \mathbf{rhs}_s^{(k)}.$$
8. Monitor the convergence for the flow subproblem:

$$\left\| \frac{\partial \mathcal{R}_a}{\partial \mathbf{W}} \frac{d\mathbf{W}}{dp_i}^{(k+1)} - \mathbf{rhs}_a^{(k)} \right\|_2 < \epsilon_a \left\| \mathbf{rhs}_a^{(0)} \right\|_2 \text{ where } \epsilon_a = 1e-6$$
9. Monitor the convergence for the structural subproblem:

$$\left\| [\mathbf{K} - \tilde{\mathbf{K}}_a] \left(\frac{d\mathbf{U}}{dp_i}^{(k+1)} - \frac{d\mathbf{U}}{dp_i}^{(k)} \right) \right\|_2 < \epsilon_s \left\| \mathbf{rhs}_s^{(0)} \right\|_2 \text{ where } \epsilon_s = 1e-6$$
10. Break the iteration only if both convergence criteria are met, else restart from 4.

The proposed method in Table [2.1](#) can be reduced to the LBGS method by a simple omission of $\tilde{\mathbf{K}}_a$. It must be clarified that, upon convergence, both methods result in the same gradient as the original linear system in Equation [2.18](#). The pertinent question, which is addressed in the following section, is whether the inclusion of $\tilde{\mathbf{K}}_a$ has a positive influence on the convergence rate. Hence, the question is related to convergence rate, not accuracy. It should also be mentioned that there are additional means of improving the convergence rate. For instance, the convergence tolerance of the flow subproblem in step 5 can be adjusted. Typically, when the flow subproblem is required to converge to a higher level of accuracy, the number of structural subiterations in step 7 is reduced, and conversely, a lower level of accuracy for the flow subproblem generally results in

additional structural subiterations. This implies a trade-off in computation time between the total number of flow subiterations and the number of mesh deformation computations, since the solution time of the structural subproblem in this work is negligible compared to the mesh deformation problem. Another possibility applicable only to the LBGS, which can have a substantial impact on the computation time, is to relax the structural response $d\mathbf{U}/dp_i$ following step 7 by applying an under-relaxation parameter β such as:

$$\frac{d\mathbf{U}^{(k+1)}}{dp_i} = \frac{d\mathbf{U}^{(k)}}{dp_i} + \beta \Delta \frac{d\tilde{\mathbf{U}}^{(k+1)}}{dp_i} \quad (2.33)$$

where $\Delta \frac{d\tilde{\mathbf{U}}^{(k+1)}}{dp_i}$ is the sensitivity increment and the solution vector in step 7 in Table 2.1. The range of the relaxation parameter is limited to $0 < \beta \leq 1$, where $\beta = 1$ corresponds to the LBGS method with no relaxation applied to the solution. The relaxation parameter can be set to a predetermined value that remains constant (static) throughout the gradient computation. Alternatively, it can be allowed to adapt during the gradient computation in compliance with a dynamic relaxation strategy. A very common dynamic relaxation strategy that determines the parameter based on the iteration history is Aitken's acceleration scheme [57]. The relaxation parameter is computed based on previous increments in accordance with Barcelos et al. [24]:

$$\beta^{(k+1)} = \beta^{(k)} \left(1 - \frac{\left(\Delta \frac{d\tilde{\mathbf{U}}^{(k+1)}}{dp_i} - \Delta \frac{d\tilde{\mathbf{U}}^{(k)}}{dp_i} \right)^T \Delta \frac{d\tilde{\mathbf{U}}^{(k+1)}}{dp_i}}{\left(\Delta \frac{d\tilde{\mathbf{U}}^{(k+1)}}{dp_i} - \Delta \frac{d\tilde{\mathbf{U}}^{(k)}}{dp_i} \right)^T \left(\Delta \frac{d\tilde{\mathbf{U}}^{(k+1)}}{dp_i} - \Delta \frac{d\tilde{\mathbf{U}}^{(k)}}{dp_i} \right)} \right) \quad (2.34)$$

Regardless of the static or dynamic relaxation strategy, an initial value for the parameter $\beta^{(0)}$ must be applied. This initial value is typically assigned based on experience and has a significant influence on the convergence rate. Applying a relaxation to the structural response, as much as it can be considered an advantage by adding stability to the convergence, is also a disadvantage. Each problem is subject to a different geometry, flight condition and structural configuration for instance and finding an ideal initial parameter beforehand can be a very difficult task. One parameter that results in very good convergence rates for a specific configuration and flight condition can suddenly become a very poor parameter by adjusting one of the aforementioned features.

2.4 CASE STUDY

A case study is designed to investigate the convergence characteristics of the proposed method in the previous section and to compare the results with the more conventional LBGS method. The case description (including aerodynamic models, structural model, design variable, computational resources and so forth) is outlined in detail in the following subsection, whereas the results are presented in Subsection [2.4.2](#)

2

2.4.1 SETUP

The ONERA M6 wing is used for this case study with 1.05 and 3.57 million cells for the Euler and the RANS problem, respectively. The Euler mesh is split into 8 blocks and executed in parallel mode on 8 processors, whereas the RANS mesh is split into 42 blocks and executed in parallel mode on 16 processors. All the computations in this work are performed on a High Performance Computing (HPC) cluster with 2.60 GHz Intel Xeon E5-2660 v3 Processors. The Mach number is fixed at 0.84 with an incidence angle of 3.06° . Hence, the aeroelastic solution is not trimmed in this particular case study. Three subcases with different dynamic pressures are examined: (i) 28.9 kPa, (ii) 57.9 kPa and (iii) 86.9 kPa. The purpose here is to investigate the performance of the proposed method when the wing is subject to a strong aeroelastic effect. This is typically the case when the aerodynamic stiffness matrix is of comparable size to the structural stiffness matrix. The Reynolds numbers based on the Mean Aerodynamic Chord (MAC), are 7.42×10^6 , 14.8×10^6 and 22.3×10^6 , respectively.

The vortex lattice model is subject to the same input parameters as its high-fidelity counterpart. It is discretized by 648 panels in total: 18 in the chordwise direction and 36 in the spanwise direction. The computation time of $\tilde{\mathbf{K}}_a$ amounts to 10 – 20 seconds and is negligible compared to the total wall time of the aerostuctural gradient computation. It is therefore omitted in the results section.

The structural model is a classic wingbox layout with 1812 DOF depicted in Figure [2.6](#). Thicknesses and sections of structural elements have been designed in a pre-processing optimization step. Small adjustments have been applied to the original model in Blondeau et al. [\[29\]](#) to accommodate for the interfacing with the vortex lattice model. Rigid body elements are attached to the ribs with load transfer nodes at $z = 0$ as can be seen in Figure [2.7](#). This allows the interpolated loads from the panel method to maintain a smooth and realistic distribution. Moreover, not every rib is modified for the said reason. There are three coupling nodes per rib. The first one is located in the middle between the wing surface leading edge and the wingbox front spar. The second one is aligned with the mid spar and the third one is located in the middle between the wing surface trailing

2. GRADIENT COMPUTATION

edge and the wingbox rear spar. The structural design variable p_i , that is used for the construction of $\partial\mathbf{K}/\partial p_i$ in this case study, is the thickness variation of the highlighted design region in Figure 2.6. This region contains four shell elements.

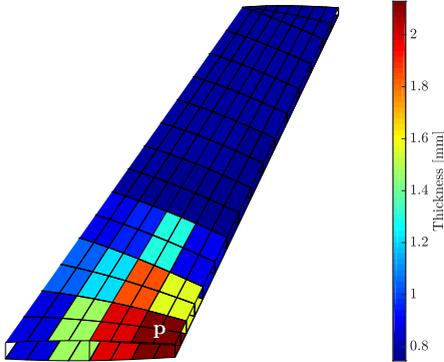


Figure 2.6: Shell thicknesses of original model

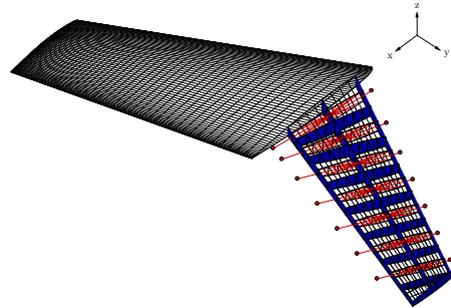


Figure 2.7: Modified model with rigid elements

In addition to comparing the proposed method, which for the sake of convenience is denoted the Multifidelity (MF) method in the remainder of this chapter, to the LBGS method, it is also benchmarked towards a pure rigid aerodynamic gradient computation at the static aeroelastic equilibrium. This corresponds to solving steps 1-5 in Table 2.1 and adjusting the convergence tolerance in step 5 to six orders of magnitude. The aero-structural gradient computation should ideally be as efficient as the rigid aerodynamic gradient computation, and this comparison will give an indication of how well the MF method is performing. Moreover, both static and dynamic under-relaxations are applied to the structural response in the LBGS method in an attempt to stabilize and accelerate the convergence. For the static relaxation strategy $\beta = 0.5, 0.7, 0.9$ and 1.0 are considered, where $\beta = 1.0$ corresponds to no relaxation. These values are kept constant throughout the gradient computation. For Aitken's adaptive strategy an initial parameter of $\beta^{(0)} = 0.5$ is applied. During the subsequent iterations, the value is adapted according to Equation 2.34.

2.4.2 RESULTS

Before the solution methods are compared, the gradients for the six subcases described in the previous subsection (three for the Euler aero-structural gradient problem and three for the RANS aero-structural gradient problem) are verified. It is important to highlight that matrix-vector products including the partial derivatives $\partial\mathcal{R}_a/\partial\mathbf{X}_a$ and $\partial\mathbf{K}/\partial p_i$ are computed by a finite difference scheme, whereas the remaining partial derivatives are obtained analytically by hand differentiating the code. Hence, the aero-structural gradients are classified as semi-analytical in

the remainder of this work. The reference aero-structural gradients in this work are obtained by forward differencing with the perturbation size $\Delta p_i = 1e-6$. The perturbation size was obtained by a separate convergence study. The quantities of interest are the lift coefficient C_L and the pressure drag coefficient C_{D_p} . The three dynamic pressures produce a static aeroelastic tip deflection of 6.8%, 12% and 15% of the semi-span. As can be seen in Table 2.2, the Euler semi-analytical approach generates gradients that are closely matching the finite difference results. The relative difference is less than 0.0024% for all three subcases. The RANS semi-analytical approach on the other hand shows larger discrepancies when compared to the finite difference results. The relative error ranges from 3.1% to 7.0% for C_L and 2.9% to 10.6% for C_{D_p} . There are multiple sources that can be attributed to this outcome. One source of error is the frozen-turbulence assumption applied to the flux Jacobian $\partial \mathcal{R}_a / \partial \mathbf{W}$. It has been shown, for instance, by Nielsen and Anderson [49] that the lift coefficient gradient can deviate for specific settings as much as 53.8% compared to a full-turbulence linearization. The comparison of linearizations in their paper was mainly performed on a 91k tetrahedra grid of the ONERA M6 wing with a Mach number of 0.3 and a Reynolds number of 5×10^6 based on the mean aerodynamic chord. In Luy et al. [51], they compare gradients obtained by the frozen-turbulence assumption and the full-turbulence linearization to a complex-step finite difference method. They show that the gradients obtained by full-turbulence linearization are as much as five orders of magnitude more accurate compared to the gradients obtained by the frozen-turbulence linearization. These two references suggest that the frozen-turbulence assumption is a likely source of error. Moreover, the perturbation size of the central difference scheme involving $\partial \mathcal{R}_a / \partial \mathbf{X}_a$ can greatly affect the outcome. A convergence study for this parameter was not conducted.

Table 2.2: Accuracy verification of the aero-structural gradients

Euler	Subcase 1		Subcase 2		Subcase 3	
	C_L	C_{D_p}	C_L	C_{D_p}	C_L	C_{D_p}
forward difference	0.971318	0.074998	1.708273	0.112197	2.122124	0.123306
semi-analytical	0.971331	0.074999	1.708290	0.112199	2.122163	0.123309
relative difference [%]	0.00130	-0.00130	0.00090	0.00180	0.00180	0.00240
RANS	C_L	C_{D_p}	C_L	C_{D_p}	C_L	C_{D_p}
forward difference	0.927958	0.062810	1.660365	0.094373	2.069982	0.100292
semi-analytical	0.899545	0.064642	1.733840	0.103102	2.214498	0.110923
relative difference [%]	-3.06190	2.91670	4.42520	9.24950	6.98150	10.6000

The convergence behavior of subcase 1 for the Euler aero-structural gradient can now be investigated. The Euclidean norm of the flow residual is plotted with respect to the total number of flow subiterations in Figure 2.8a. The LBGS ($\beta = 1.0$), i.e. the LBGS without relaxation, and the MF method have a similar rate of convergence which is lower compared to the rigid aerodynamic gradient. Both methods require the same number of flow subiterations. As the static

2. GRADIENT COMPUTATION

relaxation parameter is reduced, the number of coupling iterations increases substantially. Even Aitken's dynamic relaxation strategy, LBGS ($\beta^{(0)} = 0.5$), does not accelerate the convergence in comparison to the conventional LBGS method. This suggests that there is a strong possibility of a superior initial parameter than 0.5 for this particular case using Aitken's relaxation strategy.

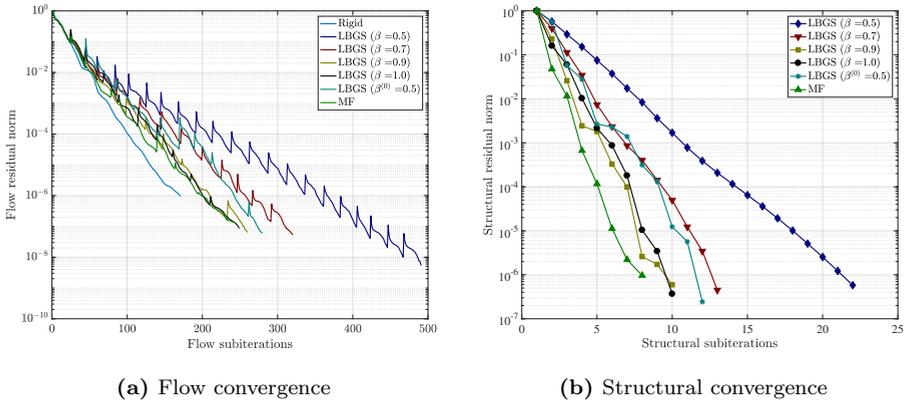


Figure 2.8: Euler case with dynamic pressure 28.9 kPa

Sharp peaks can be distinguished when the residual norm drops one order of magnitude in Figure 2.8a. This is due to the fluid-structure coupling effect when the structural subproblem is solved. In Figure 2.8b, the Euclidean norm of the structural residual is plotted with respect to the total number of structural subiterations. This plot also represents the number of mesh deformation procedures required, since each structural solution is followed by a mesh deformation. LBGS ($\beta = 1.0$) requires ten coupling iterations compared to eight for the MF method. Hence, there is only a very small computational gain to be made by the MF method due to the fewer number of mesh deformations required.

In Table 2.3, the wall time is divided in two categories that overwhelmingly dominate the total wall time: (i) the aero category and (ii) the mesh deformation category. They are represented in Table 2.1 by steps 5 and 4, respectively. The computation time of the remaining steps are negligible in comparison and are therefore omitted in this case study. As can be seen in the table, the MF method only shows a very modest 4% wall time reduction compared to the LBGS ($\beta = 1.0$) method for subcase 1, with $\beta = 1.0$ being the best parameter for this particular subcase. The small difference indicates that the effect of including the preconditioner $\tilde{\mathbf{K}}_a$ for this subcase, with a modest structural deformation, is negligible.

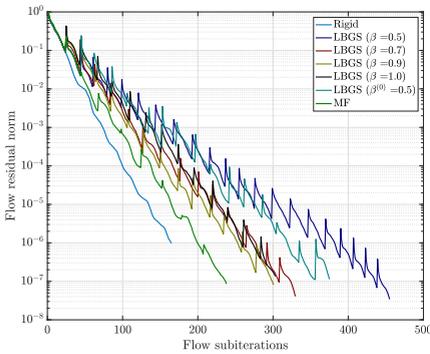
Subcase 2 for the Euler aero-structural gradient suggests that the difference between MF and LBGS becomes more pronounced as the dynamic pressure increases according to Figure 2.9a. The number of flow subiterations between MF

Table 2.3: Comparison of wall time [s]

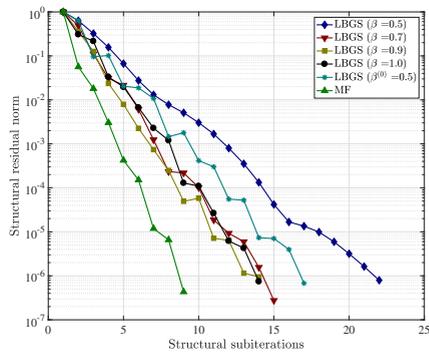
Euler	Subcase 1			Subcase 2			Subcase 3		
	aero	mesh	total	aero	mesh	total	aero	mesh	total
LBGS/ $\beta = 1.0$	670	59	729	841	85	926	1467	157	1624
LBGS/ $\beta = 0.9$	794	73	867	931	113	1044	1334	169	1503
LBGS/ $\beta = 0.7$	979	102	1081	1028	116	1144	1044	139	1183
LBGS/ $\beta = 0.5$	1517	196	1713	1410	161	1571	1381	172	1553
LBGS/ $\beta^{(0)} = 0.5$	856	101	957	1150	118	1268	1462	183	1645
MF	653	47	700	654	57	711	644	60	704
gain [%]	2.53	20.3	4.00	22.2	32.9	23.2	38.3	56.8	40.5
RANS	aero	mesh	total	aero	mesh	total	aero	mesh	total
LBGS	17840	870	18710	22288	1266	23554	25126	1907	27034
MF	18392	759	19151	18833	878	19711	22151	1007	23158
gain [%]	-5.57	12.8	-2.36	15.5	30.6	16.3	11.8	47.2	14.3

2

and LBGS ($\beta = 1.0$) differ by 65. The MF method maintains a rate of convergence similar to the rigid aerodynamic gradient, whereas the rate of convergence of the LBGS method is markedly degraded. The MF method exhibits 23% faster convergence in terms of wall time. As with the previous subcase, the reduction of the relaxation parameter does not accelerate the solution. LBGS ($\beta = 1.0$) and LBGS ($\beta = 0.9$) arguably result in the same convergence rate. The remaining relaxation parameters, including Aitken’s adaptive relaxation strategy, result in increased computation times. Figure 2.9b indicates that the number of mesh deformation operations can be reduced from 14 to 9 by the MF method.



(a) Flow convergence



(b) Structural convergence

Figure 2.9: Euler case with dynamic pressure 57.9 kPa

Subcase 3 shows that a 41% wall time reduction can be achieved by the MF method according to Table 2.3. The convergence rate of the MF method, as illustrated in Figure 2.10a, is again similar to the rigid aerodynamic gradient

as opposed to the LBGS method which becomes significantly degraded. It is interesting to note that a static relaxation parameter of 0.7 results in the fastest convergence rate for the LBGS method. This is in contradiction to the previous two subcases. As the dynamic pressure increases, and the structural deformations become substantial, the reduction of the relaxation parameter accelerates the solution of the aero-structural gradient. However, excessive relaxation, such as $\beta = 0.5$, increases the robustness of the solution but also the computation time. Again, LBGS ($\beta^{(0)} = 0.5$) results in a high computation time. This supports the premise that an initial relaxation factor of 0.5 might not be ideal for Aitken's relaxation strategy.

2

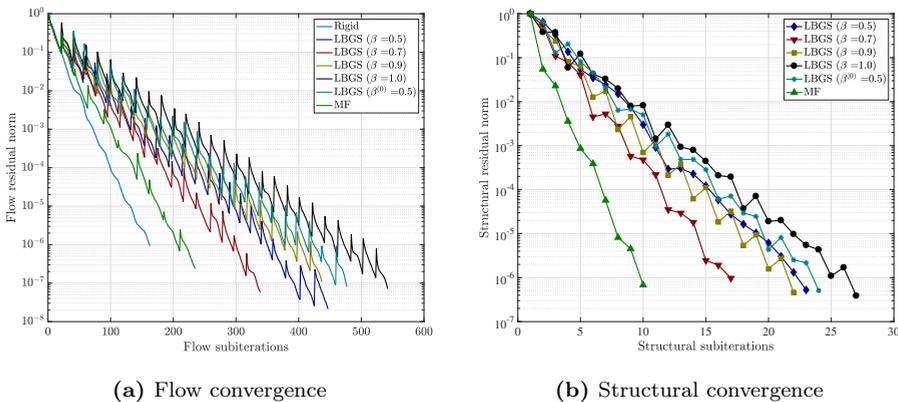


Figure 2.10: Euler case with dynamic pressure 86.9 kPa

An interesting note to make is that the wall time of the MF method is nearly identical for all three subcases. Its convergence rate seems to be almost unaffected by the change in the dynamic pressure. In contrast, the wall time of the LBGS method increases significantly with the dynamic pressure. Hence, the lack of robustness and the slow convergence rate that are the main shortcomings of the LBGS method are effectively mitigated by the MF method for this case.

If we look at the RANS aero-structural gradient, subcase 1 indicates a different convergence pattern for the rigid aerodynamic gradient, as well as the aero-structural gradients, as can be seen in Figure 2.11a. The drop in the aerodynamic residual is consistent for the Euler problem as opposed to the RANS problem which appears to change the rate after three orders of magnitude. The residual drops three orders of magnitude within roughly 250 iterations. It then requires roughly an additional 1750 iterations to drop the remaining three orders of magnitude. In addition, the convergence rate for the RANS problem is significantly lower compared to the Euler problem. This suggests that the aerodynamic linear subproblem for the RANS case has a higher condition number and thus requires an improved preconditioner for the GMRES solver. In this work, the precon-

ditioner is built using an approximate inverse (by applying several iterations of an upper-lower block relaxation) of the physics based approximate flux Jacobian $\partial\hat{\mathcal{R}}_a/\partial\mathbf{W}$. This formulation implies a fixed linear operator for the preconditioner which leads to a standard GMRES algorithm. This approach was dictated by the limitations of the ELSA version that was used in this study.

Even if the Newton-Krylov method based on restarted GMRES is more efficient than the standard BWE time marching solution method, the flexible variant of the GMRES algorithm is much more attractive in practice. A Block Flexible GMRES with Deflated Restarting [58, 59] has been recently implemented into the optimization module of the latest ELSA version. More specifically, the nested GMRES approach (i.e. the inverse of the preconditioner is obtained by calling an internal GMRES preconditioned by the approximate inverse of the $\partial\hat{\mathcal{R}}_a/\partial\mathbf{W}$) is very promising.

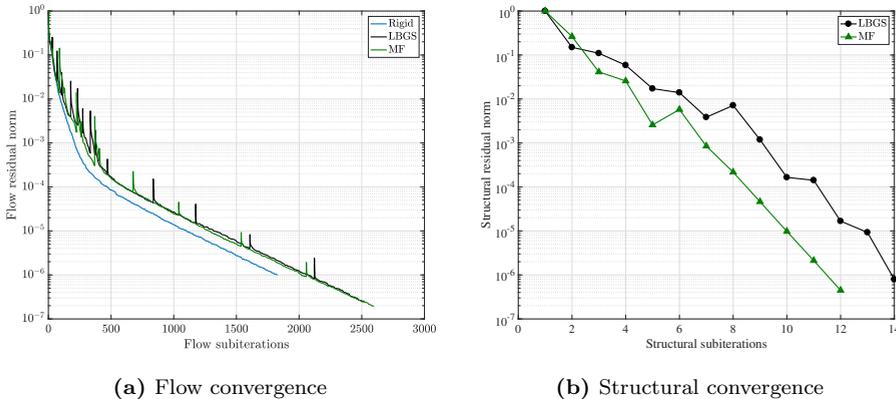
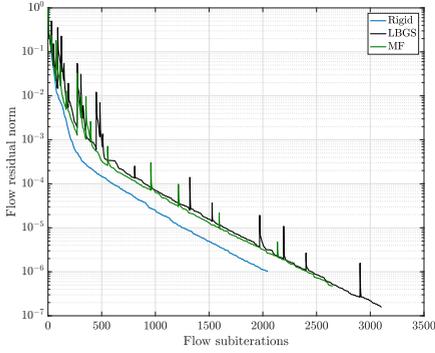


Figure 2.11: RANS case with dynamic pressure 28.9 kPa

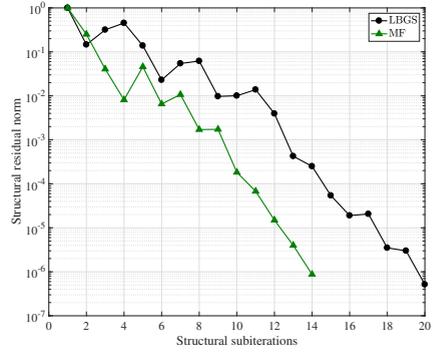
Subcase 2 for the RANS aero-structural gradient problem indicates that the number of flow subiterations can be reduced with the MF method and a wall time reduction of 16% is obtained compared to the LBGS method. A faster convergence can be observed for the MF method in [2.12b]. The total number of structural subiterations is reduced from 20 to 14. This trend extends to the third subcase where the structural subiterations are reduced from 30 to 16 using the MF method. The wall time reduction for subcase 3 is only 14%. It can be concluded that the structural subiterations are effectively kept at a minimum using the MF method, whereas the convergence rate of the flow subproblem is dominated by the dedicated aerodynamic preconditioner $\partial\hat{\mathcal{R}}_a/\partial\mathbf{W}$. For this particular case study the LU-SSOR preconditioner seems to work well for the Euler problem, but not as good for the RANS problem. It is possible that an ILU-type preconditioner based on a more accurate Jacobian approximation would accelerate the convergence.

2. GRADIENT COMPUTATION

2

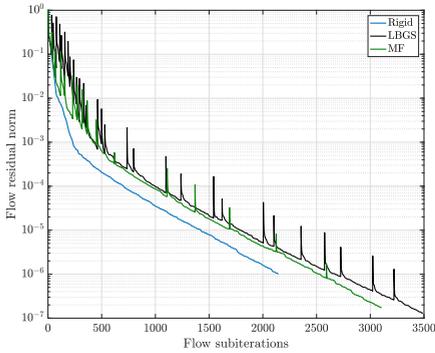


(a) Flow convergence

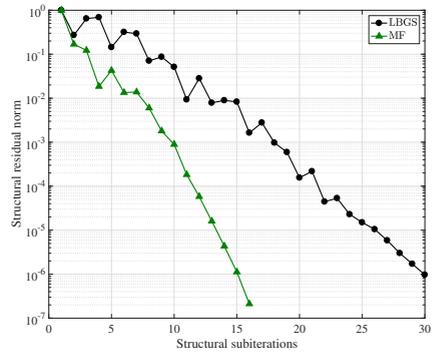


(b) Structural convergence

Figure 2.12: RANS case with dynamic pressure 57.9 kPa



(a) Flow convergence



(b) Structural convergence

Figure 2.13: RANS case with dynamic pressure 86.9 kPa

It should be mentioned that no structural relaxation is applied to LBGS method for the RANS aero-structural gradient problem. The main reason for this is to enhance the visibility and avoid clutter in the plots. Moreover, the computation time for the RANS problem is evidently dominated by the flow subproblem and the type of aerodynamic preconditioner applied. Hence, a relaxation of the structural response is not likely to result in a substantial change in the overall computation time, as opposed to the Euler problem.

2.5 CONCLUSIONS

A preconditioner based on vortex lattice aerodynamics is constructed to develop a method which stabilizes and accelerates high-fidelity aero-structural gradient computations. The efficiency of the proposed method is demonstrated by a case study using computational models of the ONERA M6 wing. The method incorporating the said preconditioner (the MF method) is proven to exhibit higher convergence rates compared to the LBGS method. The method is highly modular and only requires an additional aerodynamic stiffness matrix from an external linearized potential flow solver compared to the LBGS method. The Euler case shows promising improvements with the proposed method and the gradient computation is up to 41% more efficient compared to the LBGS method. The convergence rate of the RANS case, on the other hand, is not equally promising with a 16% reduction in computation time. The main source for this difference in the outcome of the Euler and the RANS problem is the limitation to the type of preconditioner used by the flow solver in this work, ELSA. Moreover, the gradient computations in this work are only applied to the direct approach. The proposed methods compatibility with an adjoint approach is something that requires further discussion and consideration, and is not covered in the scope of this thesis.

3

GRADIENT RECONSTRUCTION

In this chapter, a method has been developed to reconstruct high-fidelity aero-structural gradients for numerical optimization. A computationally inexpensive vortex lattice model provides a good approximation of aerodynamic loads and is used in conjunction with singular value decomposition and pseudo-inverting to reconstruct the gradients of a high-fidelity aero-structural model. The proposed method is examined on two aero-structural configurations: (i) the ONERA M6 (OM6) wing and (ii) a forward-swept wing (FSW). The aerodynamics of the first configuration is represented by a 1.05 million cell Euler model and the structure is modeled by a finite element wingbox comprising 1812 degrees of freedom. The aero-structural configuration of the OM6 was already described in Subsection [2.4.1](#). The aerodynamics of the second configuration is represented by a 643K cell Euler model and the structure is modeled by a three-dimensional Timoshenko beam. The proposed method is shown to be able to reconstruct the gradients of structural and aerodynamic responses with respect to structural sizing variables.

The remainder of the chapter is organized as follows: a brief overview on aeroelastic analysis, sensitivity analysis and the aerodynamic stiffness matrix is covered in Section [3.1](#). The gradient reconstruction method is outlined in full detail in Section [3.2](#). Two case studies based on the OM6 wing and a FSW configuration are examined in Section [3.3](#) and [3.4](#), respectively. Conclusions are discussed in Section [3.5](#).

3.1 AEROELASTIC ANALYSIS AND SENSITIVITY ANALYSIS

A crucial element in aero-structural design optimization is the solution of the governing equations of the computational models. This allows for certain quantities of interest, such as aerodynamic lift, drag and moment coefficients, as well as structural deformations and strains or stresses, to be determined. These quantities can be regarded as input variables for the optimizer in search of a better design. Let the discrete steady-state Euler flow problem be expressed in residual form as $\mathcal{R}_a(\mathbf{W}, \mathbf{X}_a) = \mathbf{0}$, where \mathbf{W} is the unknown vector of conservative variables and \mathbf{X}_a is the vector of aerodynamic grid coordinates. Moreover, let $\mathcal{R}_s(\mathbf{U}, \mathbf{X}_s) = \mathbf{K}\mathbf{U} - \mathbf{Q} = \mathbf{0}$ be the linear structural problem, where \mathbf{K} is the global structural stiffness matrix, \mathbf{X}_s are the structural grid coordinates, \mathbf{U} is the vector of structural DOF and \mathbf{Q} is the applied load vector. The static aeroelastic problem can then be defined as a concatenation of the two subproblems

$$\begin{pmatrix} \mathcal{R}_a(\mathbf{W}, \mathbf{U}) \\ \mathcal{R}_s(\mathbf{U}, \mathbf{W}) \end{pmatrix} = \mathbf{0} \quad (3.1)$$

where the constant coordinate vectors have been omitted for the sake of brevity and only unknown variables remain. The two sets of equations are coupled through aerodynamic loads acting on the structure and aerodynamic grid perturbations subject to structural deformations. In theory, both residuals should be reduced to zero to ensure accurate function values for the desired quantities of interest. In practice, however, this is rarely the case. The equilibrium conditions are typically met when the change in residuals over successive iterations is within a specified tolerance. Throughout this thesis a defect-correction approach is applied, which was outlined in full detail in Subsection 2.2. Once the static aeroelastic problem is considered solved, and the equilibrium deformations are determined, the stresses $\boldsymbol{\sigma}$ can be recovered by

$$\boldsymbol{\sigma} = \mathbf{S}\mathbf{U} \quad (3.2)$$

where \mathbf{S} is a stress-displacement matrix [60]. This matrix is assumed to be known. Theoretically, the routines that evaluate the stresses or strains for a given set of design variables would suffice for zeroth-order optimization. However, in an effort to reduce the number of function evaluations, and by extension the computational effort, the preferred option in the context of high-fidelity aeroelasticity is gradient-based optimization. As the name suggests, gradient-based optimization additionally requires sensitivity analysis to determine the total derivatives of the stresses with respect to structural sizing variables $d\boldsymbol{\sigma}/dp_i$. To obtain these gradients, the entire chain of operators for computing the stresses must be linearized. This implies that the equilibrium Equation 3.1 must be differentiated with respect

to the design variables

$$\begin{pmatrix} \frac{d\mathcal{R}_a}{dp_i} \\ \frac{d\mathcal{R}_s}{dp_i} \end{pmatrix} = \begin{pmatrix} \frac{\partial\mathcal{R}_a}{\partial p_i} \\ \frac{\partial\mathcal{R}_s}{\partial p_i} \end{pmatrix} + \begin{bmatrix} \frac{\partial\mathcal{R}_a}{\partial\mathbf{W}} & \frac{\partial\mathcal{R}_a}{\partial\mathbf{U}} \\ \frac{\partial\mathcal{R}_s}{\partial\mathbf{W}} & \frac{\partial\mathcal{R}_s}{\partial\mathbf{U}} \end{bmatrix} \begin{pmatrix} \frac{d\mathbf{W}}{dp_i} \\ \frac{d\mathbf{U}}{dp_i} \end{pmatrix} = \mathbf{0} \quad (3.3)$$

The linear system, which is commonly referred to as the Global Sensitivity Equations (GSE), was originally conceived by Sobieski [61] as a means of mathematically expressing the total sensitivity derivatives of a generic multi-disciplinary system. In this thesis, the disciplines are limited to aerodynamics and structures. To determine the total stress derivative, Equation 3.2 is differentiated with respect to the design variables using the chain rule

$$\frac{d\boldsymbol{\sigma}}{dp_i} = \mathbf{S} \frac{d\mathbf{U}}{dp_i} + \frac{\partial\mathbf{S}}{\partial p_i} \mathbf{U} \quad (3.4)$$

When the design variable is a material property such as skin thickness, and not linked to a geometric variation, the partial derivative of the stress-displacement matrix is equal to zero. Hence, the total derivative can be determined by a direct multiplication of $d\mathbf{U}/dp_i$, obtained in Equation 3.3. To simplify the expression of the GSE, $d\mathbf{W}/dp_i$ can be eliminated by Schur complement reduction

$$\underbrace{[\mathbf{K} - \mathbf{K}_a]}_{\mathbf{S}} \frac{d\mathbf{U}}{dp_i} = -\frac{\partial\mathbf{K}}{\partial p_i} \mathbf{U} \quad (3.5)$$

where the static pseudo-load $(\partial\mathbf{K}/\partial p_i)\mathbf{U}$ is equal to $\partial\mathcal{R}_s/\partial p_i$. This term is traditionally obtained by finite differencing and was derived in Subsection 2.1.2. Moreover, $\partial\mathcal{R}_a/\partial p_i$ is equal to zero since the thicknesses of the structural elements do not have a direct impact on the aerodynamic residual. The coefficient matrix \mathbf{S} of the linear system is the Schur complement of the block $\partial\mathcal{R}_a/\partial\mathbf{W}$, whereas \mathbf{K}_a is known as the aerodynamic stiffness matrix. This matrix plays a central part in this work and it describes how the aerodynamic loads change with respect to the structural DOF. Unlike the structural stiffness matrix, this term is very impractical to construct for industrially scaled problems since it requires the exact inverse of $\partial\mathcal{R}_a/\partial\mathbf{W}$ (see Equation 2.26). As a result, the linear system is often solved by iterative methods, where only matrix-vector products of the aerodynamic stiffness matrix are necessary. In this thesis, the method of choice for solving the linear system was described in full detail in Subsection 2.3. In this chapter, it is only assumed that the linear system can be solved iteratively and that it must be solved as many times as there are right-hand sides.

It should also be mentioned that the structural stiffness matrix is generally symmetric, sparse, positive definite and can be stored and factorized efficiently by sparse or regular Cholesky decomposition. The aerodynamic stiffness matrix, on the other hand, does in general not possess any of these properties. It is typically unsymmetric, dense and ill-conditioned.

3.2 GRADIENT RECONSTRUCTION

A generic optimization problem of a structural wingbox can contain several thousands of structural design variables and be subject to as many stress constraints as there are structural elements. Solving Equation 3.5 iteratively for all the design variables in a high-fidelity model framework is computationally prohibitive and therefore not a viable option. An alternative, and rather common approach, is to lump all the stress constraints into a composite function and subsequently solve the linear gradient problem by the adjoint approach. The constraint lumping can be achieved by an aggregation strategy, where the KS function [31] is one of the more popular options and can be expressed as:

3

$$KS[\mathbf{g}(\mathbf{p})] = g_{max}(\mathbf{p}) + \frac{1}{\rho} \ln \left[\sum_{j=1}^m e^{\rho(g_j(\mathbf{p}) - g_{max}(\mathbf{p}))} \right] \quad (3.6)$$

where m represents the number of constraint functions g , \mathbf{p} is the vector of structural design variables, g_{max} is the maximum value evaluated at \mathbf{p} and ρ is the KS parameter. A major drawback to this approach is the KS parameter, which must be chosen judiciously. Too high values result in sharp changes in the gradients which lead to a difficult convergence of the optimization problem, whereas too small values lead to overly conservative designs. In this chapter, an alternative approach of calculating the gradients is proposed and it is outlined in the following subsection.

3.2.1 METHODOLOGY

An assumption can be made where the aerodynamic stiffness matrix $\tilde{\mathbf{K}}_a$ (see approach in Subsection 2.1.4) can be inexpensively constructed by the VLM. On the contrary, the high-fidelity counterpart \mathbf{K}_a based on the Euler equations is assumed to be unavailable. The matrix $\tilde{\mathbf{K}}_a$ can be a very good approximation to \mathbf{K}_a if the flow is subsonic and much less so if the flow is transonic. Regardless, there is a gradient defect which constitutes the difference between the two matrices

$$\mathbf{K}_a = \tilde{\mathbf{K}}_a + \Delta\mathbf{K}_a \quad (3.7)$$

where $\Delta\mathbf{K}_a$ is the unknown defect that is to be approximated. It is well known that a diagonalizable matrix in general yields real eigenvectors. However, for the aerodynamic stiffness matrix, this property can not be guaranteed. Consequently, to avoid complex eigenvectors, the Singular Value Decomposition (SVD) is applied. The aerodynamic stiffness matrix can then be expressed as a product of

three matrices

$$\mathbf{K}_a = \mathbf{R}\mathbf{\Sigma}\mathbf{T}^T \quad (3.8)$$

where \mathbf{R} and \mathbf{T} are square matrices spanned by orthonormal basis vectors. Therefore, $\mathbf{R}^T\mathbf{R} = \mathbf{T}^T\mathbf{T} = \mathbf{I}$, where \mathbf{I} is the identity matrix. In the remainder of this chapter, \mathbf{R} and \mathbf{T} will be referred to as the left- and right-singular basis, respectively. The diagonal matrix $\mathbf{\Sigma}$ stores the singular values for the corresponding basis vectors in decreasing numerical order. Alternatively, the SVD can be considered as a sum of rank-1 matrices

$$\mathbf{K}_a = \sum_{i=1}^n \sigma_i \{\mathbf{r}_i\} \{\mathbf{t}_i\}^T \quad (3.9)$$

where σ_i is the i^{th} diagonal entry of $\mathbf{\Sigma}$ and $\{\mathbf{r}_i\}$ and $\{\mathbf{t}_i\}$ are the i^{th} columns of \mathbf{R} and \mathbf{T} , respectively. The Euclidean norm of each matrix in the summation is equal to one. This implies that the first matrices, accompanied by the highest singular values, will have the largest impact on the aerodynamic stiffness matrix. If an assumption is made where a subset of k left-singular basis vectors $\{\mathbf{r}_1, \dots, \mathbf{r}_k\}$ and the corresponding singular values $\sigma_1, \dots, \sigma_k$ are available, then \mathbf{K}_a can be conceptually approximated by first iteratively solving

$$\mathbf{K}_a \{\mathbf{t}_1, \dots, \mathbf{t}_k\} = \{\sigma_1 \mathbf{r}_1, \dots, \sigma_k \mathbf{r}_k\} \quad (3.10)$$

and subsequently multiplying the obtained right-singular basis vectors $\{\mathbf{t}_1, \dots, \mathbf{t}_k\}$ in accordance with Equation 3.9. Unfortunately, \mathbf{K}_a is an unknown quantity and so is $\{\mathbf{r}_1, \dots, \mathbf{r}_k\}$. However, $\{\tilde{\mathbf{r}}_1, \dots, \tilde{\mathbf{r}}_k\}$ can be explicitly constructed from the approximated matrix $\tilde{\mathbf{K}}_a$. Instead, the following equation can be solved

$$\mathbf{K}_a \{\hat{\mathbf{t}}_1, \dots, \hat{\mathbf{t}}_k\} = \{\tilde{\sigma}_1 \tilde{\mathbf{r}}_1, \dots, \tilde{\sigma}_k \tilde{\mathbf{r}}_k\} \quad (3.11)$$

where $\{\hat{\mathbf{t}}_1, \dots, \hat{\mathbf{t}}_k\}$ is not a subset of the right-singular basis of \mathbf{K}_a , nor is it a subset of the right-singular basis of $\tilde{\mathbf{K}}_a$. Rather, $\{\hat{\mathbf{t}}_1, \dots, \hat{\mathbf{t}}_k\}$ can be considered an approximation to a subset of the right-singular basis of \mathbf{K}_a . If Equation 3.7 is premultiplied with the obtained basis the following expression can be formulated:

$$\Delta \mathbf{K}_a \{\hat{\mathbf{t}}_1, \dots, \hat{\mathbf{t}}_k\} = \{\tilde{\sigma}_1 \tilde{\mathbf{r}}_1, \dots, \tilde{\sigma}_k \tilde{\mathbf{r}}_k\} - \tilde{\mathbf{K}}_a \{\hat{\mathbf{t}}_1, \dots, \hat{\mathbf{t}}_k\} \quad (3.12)$$

An exact matrix-vector product of the defect matrix is now available. An approximation of the defect matrix can finally be obtained by pseudo-inverting the basis that was solved for in Equation 3.11

$$\Delta \mathbf{K}_a \approx \left[\{\tilde{\sigma}_1 \tilde{\mathbf{r}}_1, \dots, \tilde{\sigma}_k \tilde{\mathbf{r}}_k\} - \tilde{\mathbf{K}}_a \{\hat{\mathbf{t}}_1, \dots, \hat{\mathbf{t}}_k\} \right] \{\hat{\mathbf{t}}_1, \dots, \hat{\mathbf{t}}_k\}^\dagger \quad (3.13)$$

whereupon \mathbf{K}_a can be reconstructed. What has been derived up to this point encapsulates the basic concept of the gradient reconstruction method. The two key features pertinent to the method are: (i) the SVD which provides a "good" subset of basis vectors and (ii) the pseudo-inverse which is used to approximate the defect matrix. However, the main concern is that the method in its current layout is not very efficient. Firstly, it is not recommended to solve Equation 3.11 to obtain $\{\widehat{\mathbf{t}}_1, \dots, \widehat{\mathbf{t}}_k\}$, as \mathbf{K}_a is generally singular or ill-conditioned at best. A more tractable approach would be to instead solve

$$\mathcal{S}\{\widehat{\mathbf{t}}_1, \dots, \widehat{\mathbf{t}}_k\} = \{\widetilde{\sigma}_1 \widetilde{\mathbf{r}}_1, \dots, \widetilde{\sigma}_k \widetilde{\mathbf{r}}_k\} \quad (3.14)$$

followed by the reconstruction of the Schur complement defect matrix:

$$\Delta \mathcal{S} \approx \left[\{\widetilde{\sigma}_1 \widetilde{\mathbf{r}}_1, \dots, \widetilde{\sigma}_k \widetilde{\mathbf{r}}_k\} - \widetilde{\mathcal{S}}\{\widehat{\mathbf{t}}_1, \dots, \widehat{\mathbf{t}}_k\} \right] \{\widehat{\mathbf{t}}_1, \dots, \widehat{\mathbf{t}}_k\}^\dagger \quad (3.15)$$

This minor modification has no influence on the final outcome. Indeed, the structural stiffness matrix is identical in the Schur complement of the Euler model as well as the VLM model, such that $\Delta \mathcal{S} \equiv \Delta \mathbf{K}_a$. However, it is important to mention that the SVD is still applied to the term that is to be corrected, which is the aerodynamic stiffness matrix.

Another important feature of the method that requires special attention is related to the construction of the left-singular basis vectors $\{\widetilde{\sigma}_1 \widetilde{\mathbf{r}}_1, \dots, \widetilde{\sigma}_k \widetilde{\mathbf{r}}_k\}$ by SVD and the construction of the pseudo-inverse. Standard routines, in state-of-the-art numerical libraries such as LAPACK [62], are typically designed to compute the SVD and the pseudo-inverse with respect to the Euclidean norm. However, in the current context, these numerical operations are applied on three-dimensional structural DOF, with three displacements and three rotations per node. The mismatch in scaling between the displacements and rotations is not accounted for by the Euclidean norm. This limitation can be accounted for if a proper norm is selected for normalization. To this end, the strain energy norm is applied.

In the following subsection, the construction of the left-singular basis with respect to the strain energy norm is explained, whereas in Subsection 3.2.3 the construction of the pseudo-inverse with respect to the strain energy norm is outlined. A solution procedure for the proposed method is outlined in Subsection 3.2.4.

3.2.2 SVD WITH RESPECT TO THE ENERGY NORM

The strain energy norm can be defined for linear structural deformations as

$$\|\mathbf{U}\|_E = \frac{1}{2} \mathbf{U}^T \mathbf{K} \mathbf{U} \quad (3.16)$$

In a similar fashion, the strain energy norm can be expressed in terms of applied loads by

$$\|\mathbf{Q}\|_E = \frac{1}{2}\mathbf{Q}^T\mathbf{K}^{-1}\mathbf{Q} \quad (3.17)$$

Note that Equation 3.16 can be recovered by replacing \mathbf{Q} with $\mathbf{K}\mathbf{U}$ in the above equation due to the self-adjoint property of the structural stiffness matrix. The matrix \mathbf{K}_a can be regarded as a mapping operator that maps vectors from one space to another. The input to the mapping are deformations and the output are loads. This line of reasoning can be conceptually thought of as

$$\mathbf{K}\mathbf{U} = \mathbf{K}_a\mathbf{U} \quad (3.18)$$

The ratio of the output to the input can be expressed in the form of a generalized Rayleigh quotient by substituting \mathbf{Q} with $\mathbf{K}_a\mathbf{U}$ in Equation 3.17

$$\sigma^2_{(\mathbf{K}_a, \mathbf{K}; \mathbf{U})} = \max_{\mathbf{U}} \frac{\|\mathbf{Q}\|_E}{\|\mathbf{U}\|_E} \equiv \max_{\mathbf{U}} \left(\frac{\mathbf{U}^T\mathbf{K}_a^T\mathbf{K}^{-1}\mathbf{K}_a\mathbf{U}}{\mathbf{U}^T\mathbf{K}\mathbf{U}} \right) \quad (3.19)$$

It has already been established that the structural stiffness matrix \mathbf{K} is symmetric, positive definite and can therefore be replaced by a Cholesky decomposition $\mathbf{K} = \mathbf{L}\mathbf{L}^T$, where \mathbf{L} is a lower triangular matrix. The generalized Rayleigh quotient can then be reformulated accordingly:

$$\left(\frac{\mathbf{U}^T\mathbf{K}_a^T\mathbf{K}^{-1}\mathbf{K}_a\mathbf{U}}{\mathbf{U}^T\mathbf{K}\mathbf{U}} \right) \rightarrow \frac{(\mathbf{U}^T\mathbf{L})\mathbf{L}^{-1}\mathbf{K}_a^T\mathbf{L}^{-T}\mathbf{L}^{-1}\mathbf{K}_a\mathbf{L}^{-T}(\mathbf{L}^T\mathbf{U})}{(\mathbf{U}^T\mathbf{L})(\mathbf{L}^T\mathbf{U})} \quad (3.20)$$

If $\mathbf{V} = \mathbf{L}^T\mathbf{U}$ and $\mathbf{Z}_a = \mathbf{L}^{-1}\mathbf{K}_a\mathbf{L}^{-T}$ are substituted above then

$$\sigma^2_{(\mathbf{Z}_a, \mathbf{V})} = \max_{\mathbf{V}} \left(\frac{\mathbf{V}^T\mathbf{Z}_a^T\mathbf{Z}_a\mathbf{V}}{\mathbf{V}^T\mathbf{V}} \right) \quad (3.21)$$

where it is shown that $\sigma_{(\mathbf{K}_a, \mathbf{K})} \equiv \sigma_{\mathbf{Z}_a}$ and the generalized Rayleigh quotient is reduced to a standard Rayleigh quotient. Hence, instead of applying SVD directly to \mathbf{K}_a , it is applied to \mathbf{Z}_a which is normalized with respect to the energy norm

$$\mathbf{Z}_a = \mathbf{R}\mathbf{\Sigma}\mathbf{T}^T \quad (3.22)$$

However, the desire is to generate a left-singular basis for \mathbf{K}_a , not \mathbf{Z}_a . This can be rectified by post-multiplying \mathbf{Z}_a from the left by \mathbf{L} and from the right by \mathbf{L}^T such that:

$$\mathbf{K}_a = (\mathbf{L}\mathbf{R})\mathbf{\Sigma}(\mathbf{L}\mathbf{T})^T \quad (3.23)$$

In conclusion, the left-singular basis that is used for solving Equation 3.14 is $\mathbf{F} = \mathbf{L}\mathbf{R}$, where \mathbf{R} is the left-singular basis from the singular value decomposition of \mathbf{Z}_a . It should be mentioned that the singular values have been omitted in \mathbf{F} . In fact, the singular values are only scaling factors that would get canceled out. This becomes evident from a closer inspection of Equation 3.12.

3.2.3 PSEUDO-INVERSE WITH RESPECT TO THE ENERGY NORM

A generic overdetermined linear system of equations can be formulated as

$$\mathbf{A}\mathbf{x} = \mathbf{b} \quad (3.24)$$

where \mathbf{A} is a coefficient matrix of size $m \times n$ with $m > n$, \mathbf{b} is of size $m \times 1$ and \mathbf{x} is of size $n \times 1$. Since the linear system is composed of more equations than unknowns, it usually has no solution. That is, \mathbf{b} is not in the column space of \mathbf{A} . Fortunately, an approximate solution can be found by the least-square approximation which is defined as the solution $\hat{\mathbf{x}}$ that minimizes the Euclidean norm of the residual

$$\mathbf{r} = \mathbf{b} - \mathbf{A}\hat{\mathbf{x}} \quad (3.25)$$

Only instead of minimizing the norm based on the Euclidean distance, a more proper norm in the current context would be based on the strain energy:

$$S = \min(\|\mathbf{r}\|_E) \quad (3.26)$$

where $\|\mathbf{r}\|_E = \mathbf{r}^T \mathbf{K} \mathbf{r}$ is the strain energy of \mathbf{r} and S is the norm. If the residual \mathbf{r} is substituted by $\mathbf{b} - \mathbf{A}\mathbf{x}$, the above equation can be expanded as

$$S = \min(\mathbf{b}^T \mathbf{K} \mathbf{b} - (\mathbf{b}^T \mathbf{K} \mathbf{A})\mathbf{x} - \mathbf{x}^T (\mathbf{A}^T \mathbf{K} \mathbf{b}) + \mathbf{x}^T (\mathbf{A}^T \mathbf{K} \mathbf{A})\mathbf{x}) \quad (3.27)$$

In order to find the minimum, the derivative of S with respect to \mathbf{x} is equated to zero

$$\frac{\partial S}{\partial \mathbf{x}} = -2(\mathbf{A}^T \mathbf{K})\mathbf{b} + 2(\mathbf{A}^T \mathbf{K} \mathbf{A})\mathbf{x} = \mathbf{0} \quad (3.28)$$

By solving for \mathbf{x} , the least squares solution $\hat{\mathbf{x}}$ is obtained

$$\hat{\mathbf{x}} = \underbrace{(\mathbf{A}^T \mathbf{K} \mathbf{A})^{-1} (\mathbf{A}^T \mathbf{K})}_{\mathbf{A}^\dagger} \mathbf{b} \quad (3.29)$$

where \mathbf{A}^\dagger is the pseudo-inverse based on the strain energy norm. This expression can be further simplified by substituting \mathbf{K} with a Cholesky decomposition $\mathbf{L}\mathbf{L}^T$. The pseudo-inverse then becomes

$$\mathbf{A}^\dagger = (\mathbf{A}^T \mathbf{L}\mathbf{L}^T \mathbf{A})^{-1} \mathbf{A}^T \mathbf{L}\mathbf{L}^T \quad (3.30)$$

which can be further simplified as

$$\mathbf{A}^\dagger = (\mathbf{L}^T \mathbf{A})^\dagger \mathbf{L}^T \quad (3.31)$$

In conclusion, for any vector or overdetermined coefficient matrix \mathbf{A} , its corresponding pseudo-inverse with respect to the strain energy norm will be computed according to Equation [3.31](#)

3.2.4 SOLUTION PROCEDURE

The proposed method of reconstructing \mathcal{S} is outlined in Table [3.1](#)

Table 3.1: Proposed method

-
1. Given \mathbf{K} , compute the lower matrix of the Cholesky decomposition: $\mathbf{L} = chol(\mathbf{K})$.
 2. Given $\tilde{\mathbf{K}}_a$ and \mathbf{L} , compute $\tilde{\mathbf{Z}}_a = \mathbf{L}^{-1} \tilde{\mathbf{K}}_a \mathbf{L}^{-T}$.
 3. Obtain a left-singular basis $\tilde{\mathbf{R}}$ by SVD: $[\tilde{\mathbf{R}}, \tilde{\mathbf{\Sigma}}, \tilde{\mathbf{T}}] = svd(\tilde{\mathbf{Z}}_a, n)$, where n is the desired number of singular vectors.
 4. Compute $\tilde{\mathbf{F}} = \mathbf{L}\tilde{\mathbf{R}}$.
 5. Solve the linear systems $\mathcal{S}\hat{\mathbf{T}} = \tilde{\mathbf{F}}$ by an iterative scheme.
 6. Compute the normalized pseudo-inverse of $\hat{\mathbf{T}}$ by $\hat{\mathbf{T}}^\dagger = (\mathbf{L}^T \hat{\mathbf{T}})^\dagger \mathbf{L}^T$.
 7. Approximate the Schur complement gradient defect by $\Delta \mathcal{S} = (\tilde{\mathbf{F}} - \tilde{\mathcal{S}}\hat{\mathbf{T}})\hat{\mathbf{T}}^\dagger$.
 8. Reconstruct the high-fidelity Schur complement: $\mathcal{S} = \tilde{\mathcal{S}} + \Delta \mathcal{S}$.
-

3.3 CASE STUDY I - ONERA M6

In this case study, the OM6 wing is investigated. A 1.05 million cell Euler mesh split into 8 blocks with perfect load balancing is executed in parallel on 8 processors. The surface grid of the structured Euler mesh is depicted in Figure [3.1](#). A backward-Euler (BWE) scheme is used to advance the flow solution to a steady-state and the upwind Roe scheme is selected for the spatial discretization. A MUSCL interpolation is applied for second-order spatial accuracy augmented by a Van Albada limiting function (see Hirsch [39](#) for instance). The Mach number

3. GRADIENT RECONSTRUCTION

is 0.84 with a fixed incidence angle of 3.06° . Hence, the aeroelastic solution is not trimmed in this particular case. The dynamic pressure is 86.9kPa, resulting in a Reynolds number of 22.3×10^6 based on the mean aerodynamic chord. The settings in this case study are identical to the third subcase in Section 2.4

The VLM model maintains the same flow parameters as the Euler model. It is discretized by 648 panels in total: 18 in the chordwise direction and 36 in the spanwise direction. This discretization was found to yield the best trade-off between accuracy and computational time. The VLM mesh is depicted in Figure 3.2. The panels are distributed across the camber surface.

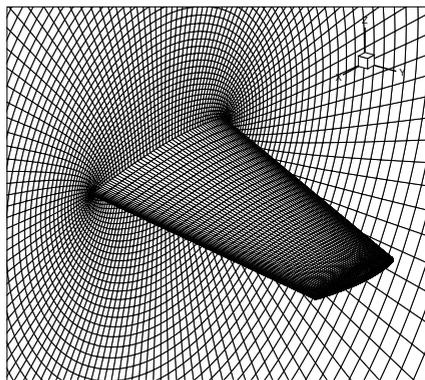


Figure 3.1: OM6 Euler surface grid

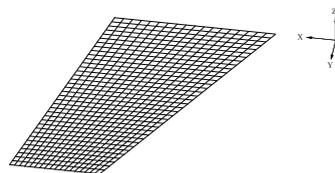


Figure 3.2: OM6 VLM surface grid

The main purpose of this case study is to examine the accuracy of the reconstructed gradients by the proposed method as the left-singular basis in step 4 of Table 3.1 is enriched. To this end, exact (reference) gradients for a subset of design variables are computed as proposed in Table 2.1, which have been verified by forward differencing. A total of 16 design regions, 8 on the top skin and 8 on the bottom skin, are selected as can be seen in Figures 3.3 and 3.4, respectively.

The static equilibrium is converged to near machine precision, with a resulting 15% out-of-plane tip deflection relative to the semi-span. The converged aerodynamic pressure distribution at equilibrium is compared in Figure 3.5. Due to the differences in discretization of the Euler and VLM models, the closest to a proper comparison is to compute the pressure difference ΔC_p of the upper and lower surface. Despite the strong aeroelastic deflection, the "delta" shock is clearly visible. This type of nonlinearity can not be modeled by the VLM.

For the sake of clarity, the transferred aeroloads are depicted in Figure 3.6. In general, there does not seem to be much difference for the leading and trailing coupling nodes, except the leading tip node. This node, along with the center coupling nodes, exhibits much higher amplitudes for the Euler model. The shock generates a strong pressure discontinuity at the transition. From a closer inspec-

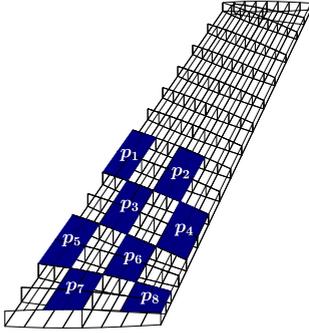


Figure 3.3: Design regions 1 to 8 on the top skin

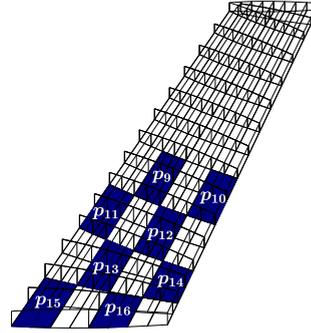


Figure 3.4: Design regions 9 to 16 on the bottom skin

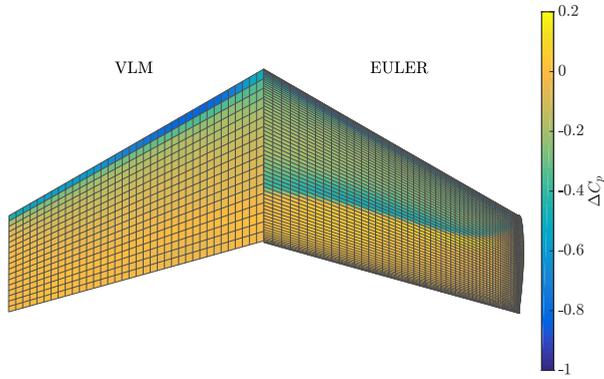


Figure 3.5: Pressure difference at camber surface

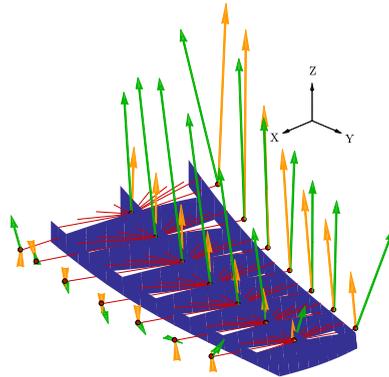


Figure 3.6: Aerodynamic loads at static aeroelastic equilibrium

3. GRADIENT RECONSTRUCTION

tion of the shock position in Figure 3.5, the results in Figure 3.6 become clear.

As confirmed by Equation 3.9, the basis vectors corresponding to the highest singular values contribute most in the reconstruction process. The number of left-singular basis vectors corresponding to non-zero singular values, in step 3 of Table 3.1 for this particular case equals 35. An interesting feature is that the number of non-zero singular values that are generated by the SVD is in proportion to the number of coupling nodes. In comparison to $\tilde{\mathbf{Z}}_a$, the reduced structural stiffness matrix \mathbf{K} generates 552 non-zero singular values, which is the size of the matrix itself. The metric that is used to determine the residual norm (i.e. the error) of the reference gradients and the reconstructed gradients is based on the strain energy

$$\|res\|_E = \frac{1}{2} \left(\frac{d\mathbf{U}^{ref}}{dp_i} - \frac{d\mathbf{U}^{rec}}{dp_i} \right)^T \mathbf{K} \left(\frac{d\mathbf{U}^{ref}}{dp_i} - \frac{d\mathbf{U}^{rec}}{dp_i} \right) \quad (3.32)$$

where the subscript i specifies the current design region, the superscript *ref* stands for the reference gradients and *rec* stands for reconstructed gradients. As can be observed in Figure 3.7, the residual norm drops considerably for the entire subset of design regions as the left-singular basis is enriched. With only 12 basis vectors applied, it is interesting to note that reconstructed gradients of design regions 10 and 11 have marginally higher residuals than their respective VLM gradients. Unfortunately, this indicates that the error is not reducing monotonically as the basis is enriched. Moreover, Figure 3.7 implies that the residual is not fully attenuated as the basis reaches the limit of 35 left-singular vectors. This makes sense since the basis vectors generated by the VLM model only approximates the corresponding inaccessible basis vectors from the Euler model. As the basis limit is reached the maximum residual drop is measured for design region 8 with 3 orders of magnitude, whereas the minimum drop is measured for design region 2 with roughly 2 orders of magnitude.

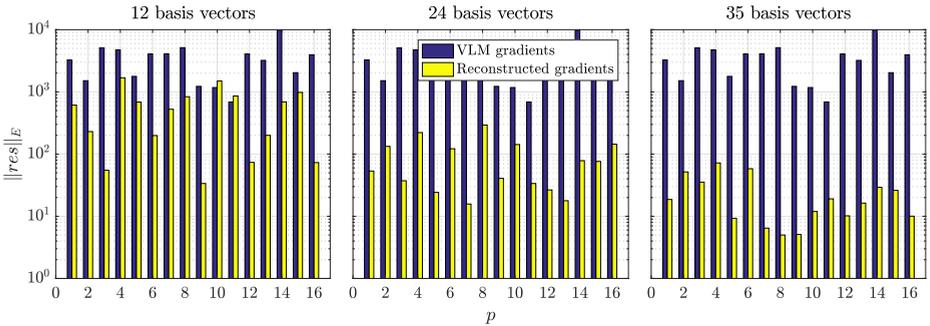


Figure 3.7: Residual error

Only structural responses have been considered thus far. However, it should be

3.4. CASE STUDY II - FORWARD SWEEP WING

mentioned that gradients of aerodynamic responses with respect to structural parameters could be computed as well. As a demonstration, the lift coefficient gradients dC_L/dp_i are estimated. This can be accomplished in a three-step procedure:

1. Approximate the aeroload gradients: $\frac{d\mathbf{Q}}{dp_i} \approx (\tilde{\mathbf{K}}_a + \Delta\mathbf{K}_a)\frac{d\mathbf{U}^{rec}}{dp_i}$
2. Compute the lift gradient: $\frac{dL}{dp_i} = \sum_{m=1} \left[\frac{dQ_m^z}{dp_i} \cos(\alpha) - \frac{dQ_m^x}{dp_i} \sin(\alpha) \right]$, where α is the incidence angle and m the coupling node index.
3. Compute the lift coefficient gradient: $\frac{dC_L}{dp_i} = \frac{dL}{dp_i}/(qS)$, where q is the dynamic pressure and S is the surface area.

The lift coefficient gradients are depicted in Figure 3.8.

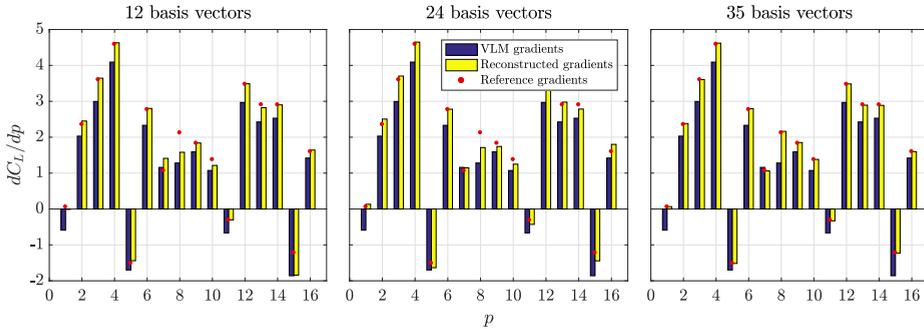


Figure 3.8: Lift coefficient gradient

It is interesting to note that the trend of the reference gradients is captured quite well by the VLM, albeit the absolute values are completely off for most design regions. The outcome is similar to that of the structural response gradients. As more basis vectors are stacked up, the defect term $\Delta\mathbf{K}_a$ improves and this is reflected by the computed gradients. Figure 3.8 shows a near-perfect match between the reference gradients and the reconstructed gradients for a vector basis of 35. It is important to emphasize that there are still very small discrepancies that can not be fully attenuated, as was highlighted in Figure 3.7. Consequently, if very accurate gradients of lift or drag coefficients are required, the most suitable method would be an exact computation of the gradients by the adjoint approach.

3.4 CASE STUDY II - FORWARD SWEEP WING

In this case study, a FSW is investigated. The main geometric characteristics are displayed in Table 3.2. The wing is equipped with a OALE10 airfoil profile at

3. GRADIENT RECONSTRUCTION

the root which is linearly interpolated to match a supercritical OAT15A airfoil profile at 25% of the semispan. The remaining span section of the wing keeps a constant OAT15A airfoil profile. The airfoil profiles are depicted in Figure 3.9. The Euler finite-volume model is discretized with 643K hexahedra split into 36 blocks, which was delivered by ONERA as part of a collaboration. The analysis and sensitivity analysis are performed in parallel mode on 20 processors with a 4.96% load imbalance. The temporal and spatial discretization schemes are identical to the previous case study.

Table 3.2: Main geometric features of the FSW

Semi-span [m]	19.10
Quarter chord sweep [deg]	18.15
Wing surface [m ²]	160
Aspect ratio [-]	9.45
Taper ratio [-]	0.265

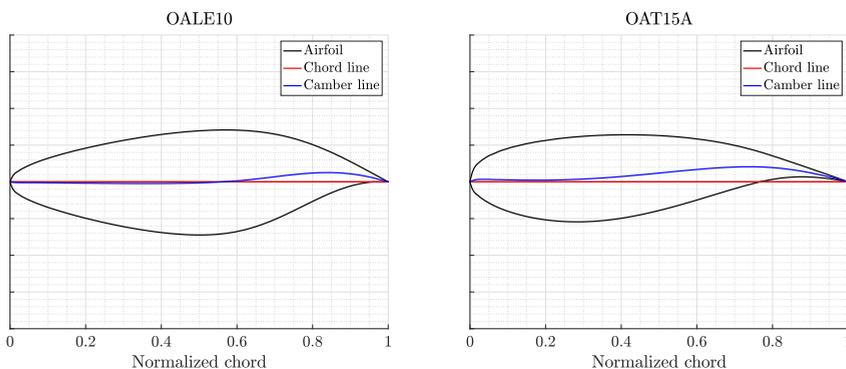


Figure 3.9: FSW airfoils. Thicknesses are not proportional to chord lengths.

For this case study, the TU Delft in-house aeroelastic design framework PROTEUS [36, 35] is employed as a tool for parameterizing and modeling the structure. PROTEUS has a built-in unsteady VLM and performs closely coupled static and dynamic aeroelastic analysis and fully analytical sensitivity analysis. The stiffness properties of the wingbox model are described by lamination parameters and are translated to a geometrically nonlinear 3D Timoshenko beam by a cross-sectional modeler. The analysis and sensitivity analysis are performed on the beam and the wingbox panel responses and gradients are retrieved in a post-treatment sequence by the cross-sectional modeler. The wingbox configuration is depicted in Figure 3.11. It consists of a front spar 15% aft of the leading edge, a rear spar 60% aft of the leading edge and a total of 25 ribs equally spaced in the span wise direction. The top and the bottom skins of the wingbox have a total of 24 panels each: 2 in the chordwise direction and 12 in the spanwise direction. The spars have 12 panels each in the spanwise direction, whereas the ribs have a constant

stiffness and thickness. This brings about a total of 72 panels and 648 design variables (8 lamination parameters and one thickness parameter per panel). The rib properties remain constant and are thus not considered as design variables. The beam is located in the center of the wingbox and is modeled by 12 elements subject to a clamped boundary condition at the root. The VLM model is depicted in Figure 3.10 with an equally spaced grid of 11 panels in the chordwise direction and 47 panels in the spanwise direction. This panel distribution was found to yield the best trade-off between accurate loads and computational time.

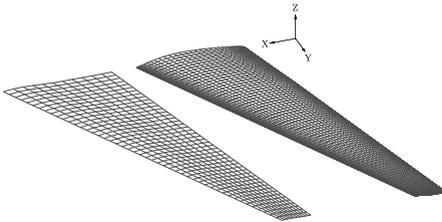


Figure 3.10: VLM panels (left) and Euler surface grid (right)

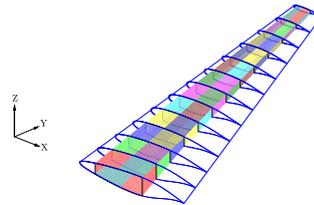


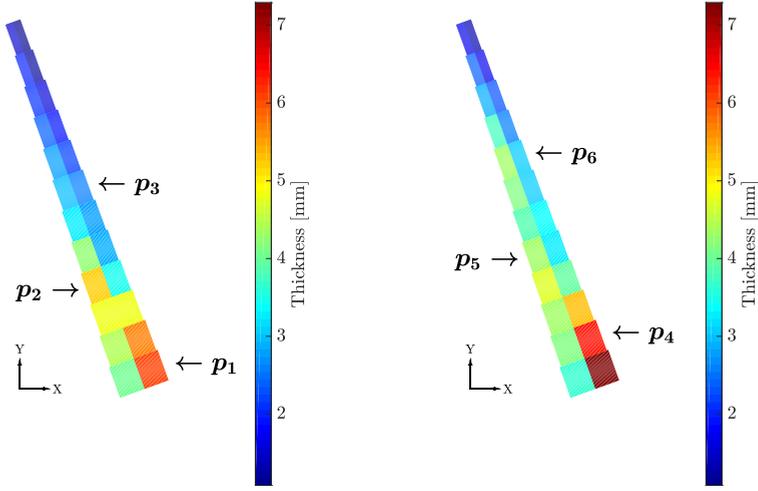
Figure 3.11: The structural wingbox

Only thickness design variables on the top and bottom skins are considered. Directional stiffness design variables, i.e. lamination parameters, are omitted for the sake of clarity in the results section. A pre-designed thickness distribution is used for this case study and is depicted in Figure 3.12. The selected design regions that are used to compute the reference gradients are highlighted for the top and bottom skin in Figures 3.12a and 3.12b. A total of 6 regions are studied.

Two significant changes are made in the current case study. Firstly, to gain a deeper understanding of the applicability of the VLM aerodynamic stiffness matrix as an approximation, three subcases are constructed with varying Mach numbers at a fixed dynamic pressure. By fixing the dynamic pressure, the local variation in static pressure becomes an isolated effect of the Mach number. This will provide more insight in how susceptible the approximation is, and in extension how well the method performs, in the presence of shocks of varying strengths. Secondly, the static aeroelastic Equations 3.1 are extended to account for trimming. In the following subsection a brief overview is provided on the necessary modifications of the PROTEUS framework to accommodate for an external CFD model. The results are then analyzed in Subsection 3.4.2.

3.4.1 EXTENSION TO TRIM ANALYSIS AND SENSITIVITY ANALYSIS

In general, the static aeroelastic equilibrium is not solved for a fixed incidence angle, but by trimming the aircraft for a certain flight condition, resulting in an additional equilibrium equation where the total lift equals the aircraft weight



(a) Design regions 1-3 on the top skin

(b) Design regions 4-6 on the bottom skin

Figure 3.12: Skin thicknesses of the structural wingbox. Flow direction is in positive x-direction.

and an additional degree of freedom in the incidence angle. The governing equations of a trimmed static aeroelastic solution in PROTEUS can be expressed as a concatenation of a structural residual \mathcal{R}_s and a trim residual \mathcal{R}_t :

$$\begin{pmatrix} \mathcal{R}_s = \mathbf{Q}_s - \tilde{\mathbf{Q}}_a - \mathbf{Q}_{ext} \\ \mathcal{R}_t = W - \tilde{L} \end{pmatrix} = 0 \quad (3.33)$$

where \mathbf{Q}_s is the internal structural load vector, $\tilde{\mathbf{Q}}_a$ is the aerodynamic load vector, \mathbf{Q}_{ext} is an external load vector associated to gravity, W is the total weight of the aircraft and \tilde{L} is the aerodynamic lift. The tilde notation is used to distinguish terms that are obtained by the VLM. Linearizing the coupled Equations 3.33 with respect to the structural DOF and the incidence angle yields

$$\begin{bmatrix} \mathbf{K} - \mathbf{K}_{ext} - \tilde{\mathbf{K}}_a & - \left(\frac{\tilde{\mathbf{Q}}_a}{\partial \alpha} + \frac{\mathbf{Q}_{ext}}{\partial \alpha} \right) \\ - \frac{\partial \tilde{L}}{\partial \mathbf{U}} & - \frac{\partial \tilde{L}}{\partial \alpha} \end{bmatrix} \begin{pmatrix} \Delta \mathbf{U} \\ \Delta \alpha \end{pmatrix} = - \begin{pmatrix} \mathcal{R}_s \\ \mathcal{R}_t \end{pmatrix} \quad (3.34)$$

where $\Delta \mathbf{U}$ and $\Delta \alpha$ are the incremental structural DOF and the incidence angle, respectively. The sensitivity analysis is commenced when the solution has converged to a state of equilibrium. The coupled Equations 3.33 are linearized with

respect to the Timoshenko beam cross-sectional stiffness matrix \mathbf{C}

$$\begin{bmatrix} \mathbf{K} - \mathbf{K}_{ext} - \tilde{\mathbf{K}}_a & - \left(\frac{\tilde{\mathbf{Q}}_a}{\partial\alpha} + \frac{\mathbf{Q}_{ext}}{\partial\alpha} \right) \\ - \frac{\partial\tilde{L}}{\partial\alpha} & - \frac{\partial\tilde{L}}{\partial\alpha} \end{bmatrix} \begin{pmatrix} \frac{d\mathbf{U}}{d\mathbf{C}} \\ \frac{d\alpha}{d\mathbf{C}} \end{pmatrix} = - \begin{pmatrix} \frac{\partial\mathcal{R}_s}{\partial\mathbf{C}} \\ \frac{\partial\mathcal{R}_t}{\partial\mathbf{C}} \end{pmatrix} \quad (3.35)$$

and the thicknesses \mathbf{t} of the design regions which are also associated to the gravity loads

$$\begin{bmatrix} \mathbf{K} - \mathbf{K}_{ext} - \tilde{\mathbf{K}}_a & - \left(\frac{\tilde{\mathbf{Q}}_a}{\partial\alpha} + \frac{\mathbf{Q}_{ext}}{\partial\alpha} \right) \\ - \frac{\partial\tilde{L}}{\partial\alpha} & - \frac{\partial\tilde{L}}{\partial\alpha} \end{bmatrix} \begin{pmatrix} \frac{d\mathbf{U}}{d\mathbf{t}} \\ \frac{d\alpha}{d\mathbf{t}} \end{pmatrix} = - \begin{pmatrix} \frac{\partial\mathcal{R}_s}{\partial\mathbf{t}} \\ \frac{\partial\mathcal{R}_t}{\partial\mathbf{t}} \end{pmatrix} \quad (3.36)$$

The solution vectors of Equations 3.35 and 3.36 are post-processed by the cross-sectional modeler in order to evaluate the total panel derivatives of various responses, such as strength or buckling, with respect to the design variables. It is important to note that while the formulation of the coefficient matrix is consistent for both analysis and sensitivity analysis, the right-hand sides are not. The pseudo-loads, i.e. the right-hand sides of Equations 3.35 and 3.36, are not influenced by the VLM when the problem is limited to structural sizing and are provided analytically in a post-analysis procedure.

The first challenge to successfully integrating an Euler flow model in the PROTEUS framework, is to ensure that the structural response is consistent with aerodynamic loads generated by the Euler model. To this end, a defect-correction approach is employed, where the loads $\tilde{\mathbf{Q}}_a$ and the lift \tilde{L} on the right-hand side of Equation 3.34 are corrected by their corresponding defects: $\Delta\mathbf{Q}_a$ and ΔL . The defects constitute the difference in aerodynamic loads between the Euler model and the VLM. Necessary inputs for the Euler flow analysis are: (i) the incidence angle, (ii) the structural DOF which are used for the mesh deformation, (iii) the Mach number and (iv) the dynamic pressure. As the solution eventually converges to the equilibrium state, it will no longer be governed by the VLM model, but by the Euler model. It is important to note that the coefficient matrix only influences the convergence rate and not the converged state. It can therefore be computationally beneficial to only compute the coefficient matrix for the first Newton iteration and keep it fixed for the subsequent iterations.

The second challenge is to rectify the solution gradients in Equations 3.35 and 3.36, such that they are consistent with the Euler model. Since the gradient problem is linear, in contrast to the static aeroelastic analysis which is geometrically nonlinear, the coefficient matrix must include the corresponding Euler flow partial derivatives. That is, the four partial derivatives distinguished by the tilde notation must be updated accordingly. The incidence angle derivative in Equation

[3.35](#) can be eliminated by a Schur complement reduction:

$$\underbrace{(\mathbf{K} - \mathbf{K}_{ext} - \mathbf{K}_a - \mathbf{K}_t)}_{\mathbf{S}} \frac{d\mathbf{U}}{d\mathbf{C}} = -\frac{\partial \mathcal{R}_s}{\partial \mathbf{C}} + \left(\frac{\mathbf{Q}_a}{\partial \alpha} + \frac{\mathbf{Q}_{ext}}{\partial \alpha} \right) \left(-\frac{\partial L}{\partial \alpha} \right)^{-1} \left(-\frac{\partial \mathcal{R}_t}{\partial \mathbf{C}} \right) \quad (3.37)$$

where \mathbf{S} is the Schur complement of the lower diagonal block in the coefficient matrix and \mathbf{K}_t is an additional stiffness term associated to the trim residual \mathcal{R}_t :

$$\mathbf{K}_t \equiv - \left(\frac{\mathbf{Q}_a}{\partial \alpha} + \frac{\mathbf{Q}_{ext}}{\partial \alpha} \right) \left(-\frac{\partial L}{\partial \alpha} \right)^{-1} \left(-\frac{\partial L}{\partial \mathbf{U}} \right) \quad (3.38)$$

3

The Schur complement formulation in Equation [3.37](#) also holds for Equation [3.36](#) but is omitted here for the sake of brevity. The procedure to obtain the matrix-vector product of the aerodynamic stiffness matrix \mathbf{K}_a was discussed in detail in Section [2.3](#) and is not repeated here. The partial derivatives $\partial \mathbf{Q}_a / \partial \alpha$ and $\partial L / \partial \alpha$ are straightforward to evaluate by forward differencing. A small perturbation $\Delta \alpha$ is applied to the trim angle, whereby the flow solution is restarted and converged to at least the same level of accuracy as the previously converged state. $\Delta \alpha = 1e-3$ was found to yield accurate gradients for this particular case. These partial derivatives require only one additional flow solution and are thus inexpensive to evaluate. The fourth partial derivative $\partial L / \partial \mathbf{U}$, or rather the matrix-vector product $(\partial L / \partial \mathbf{U})(d\mathbf{U} / d\mathbf{C})$, can be obtained directly from the construction of the matrix-vector product of \mathbf{K}_a . The procedure is identical to the construction of the lift gradient explained in the previous case study. Alternatively, $\partial L / \partial \mathbf{U}$ could be constructed explicitly and efficiently by the adjoint method. In extension, this would allow for an explicit construction of \mathbf{K}_t , i.e. \mathbf{K}_t would be constructed as a matrix and not a matrix-vector product. However, this approach would incur additional computational costs and is therefore not implemented. Once the routines for the partial derivatives are established and the iterative scheme has converged to a solution $d\mathbf{U} / d\mathbf{C}$ in Equation [3.37](#), the total incidence angle gradient $d\alpha / d\mathbf{C}$ can, if needed, be recovered by:

$$\frac{d\alpha}{d\mathbf{C}} = \left(-\frac{\partial L}{\partial \alpha} \right)^{-1} \left(-\frac{\partial \mathcal{R}_t}{\partial \mathbf{C}} + \frac{\partial L}{\partial \mathbf{U}} \frac{d\mathbf{U}}{d\mathbf{C}} \right) \quad (3.39)$$

It should be mentioned that the detailed procedure explained above to obtain $d\mathbf{U} / d\mathbf{C}$ and $d\alpha / d\mathbf{C}$ must be replicated for the thickness variable \mathbf{t} to acquire $d\mathbf{U} / d\mathbf{t}$ and $d\alpha / d\mathbf{t}$. It is only omitted here for the sake of brevity.

3.4.2 RESULTS

The Mach numbers of the three subcases are 0.72, 0.81 and 0.87, respectively. The dynamic pressure is kept constant at 22.7kPa and the Reynolds numbers of the three subcases are 33.6×10^6 , 39.0×10^6 and 43.6×10^6 with respect to the mean aerodynamic chord. The total aircraft weight W includes (i) structural masses, (ii) non-structural masses, (iii) fuel masses and (iv) remaining masses such as fuselage and empennage masses. For this particular case, the weight is estimated to 49,500 kg. The static aeroelastic solutions for the three subcases are converged to near-machine precision and are depicted in a ΔC_p plot in Figure 3.13. A comparison is made between the VLM on the left and the Euler flow on the right. It can be seen that the VLM pressure for the first subcase is capturing the general features of the Euler flow pressure very well. The Euler flow solution for the second subcase seems to generate a moderate shock at roughly 75% aft of the leading edge. As the figure shows, no shock is present in the VLM solution and it can therefore not be fully relied upon in transonic flow. At Mach 0.87, the shock intensifies in strength and translates much closer to the trailing edge.

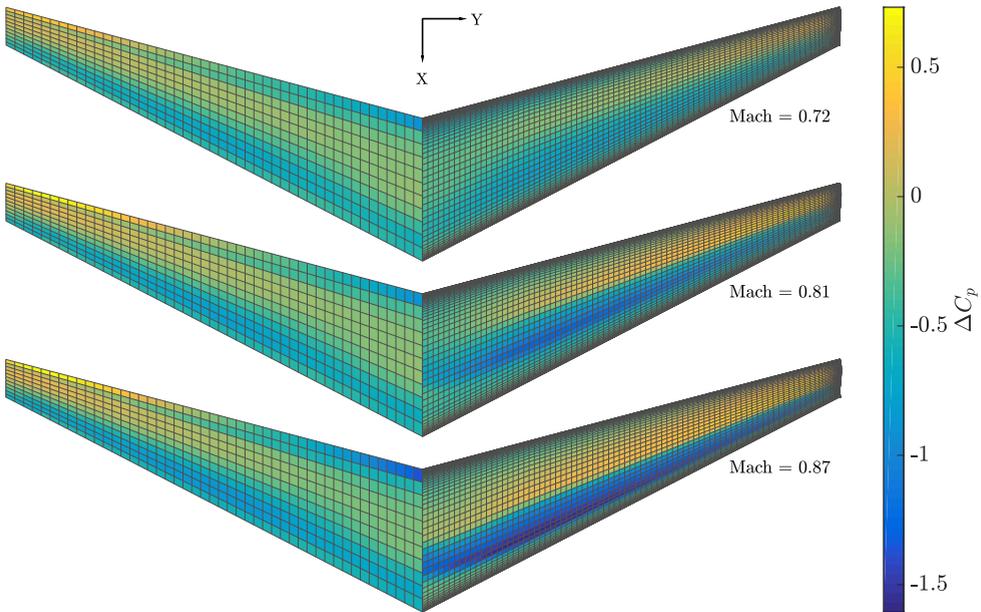


Figure 3.13: ΔC_p plot comparing VLM (left) to Euler (right) for subcase 1 (top) to subcase 3 (bottom). Flow direction is in positive x-direction.

To keep the computational cost at a minimum, a linear beam model was selected for this case study. The number of relevant singular values of $\tilde{\mathbf{Z}}_a$ is 61. Singular values above 61 drop below 10^{-15} , as can be seen in Figure 3.14, rendering the basis vectors of these singular values irrelevant. All the beam nodes in this

3. GRADIENT RECONSTRUCTION

case study are applied as coupling nodes for the load and displacement transfer. Moreover, it must be stressed that the approximation matrix, used to generate $\tilde{\mathbf{Z}}_a$, is no longer simply $\tilde{\mathbf{K}}_a$. A closer look at Equation 3.37 indicates that the proper matrix to be used for approximating the left-singular basis is $(\tilde{\mathbf{K}}_a + \tilde{\mathbf{K}}_t)$ since $\tilde{\mathbf{K}}_t$ is directly linked to the aerodynamics. Apart from step 2 in Table 3.1, the remaining steps remain exactly the same.

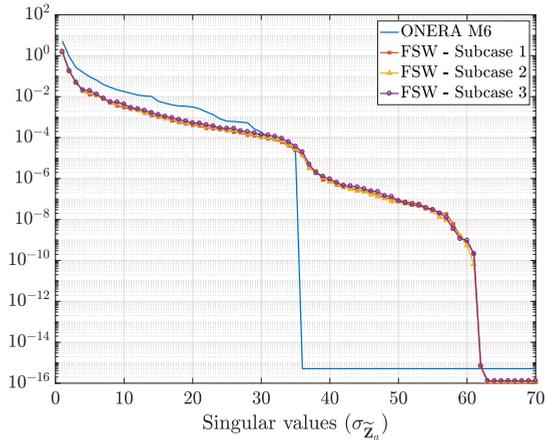


Figure 3.14: Singular values for both case studies

Subcase 1, with no shock present, indicates a very good gradient reconstruction for the selected subset of design regions according to Figure 3.15. The initial residual norm, i.e. without a gradient reconstruction, indicates the difference between the Euler gradients and the VLM gradients. The largest initial residual norm is roughly three orders of magnitude for design regions one and three and the smallest is two orders of magnitude for design region five. The largest final residual norm (i.e. when the left-singular basis is fully enriched), is roughly two orders of magnitude for design region one and the smallest is 0.4 corresponding to design region six. The design region with the best improvement over the entire span of basis vectors is design region three, which showed a residual reduction of three orders of magnitude and the least improvement is two orders of magnitude observed for design region five. It is evident, more so than the previous case study, that all the design regions are subject to strong variations and are not reducing monotonically. Nevertheless, the overall trend is clear and the residuals are reducing as the left-singular basis is enriched.

Subcase 2, where a moderate shock is present, indicates a higher initial residual for design regions 3 and 6 for instance, compared to the first subcase as can be seen in Figure 3.16. There seems to be a strong initial residual reduction for design regions 1, 3, 4 and 6. However, after roughly 20 basis vectors, the residual reduction is almost leveled out and the remaining improvement is marginal. The

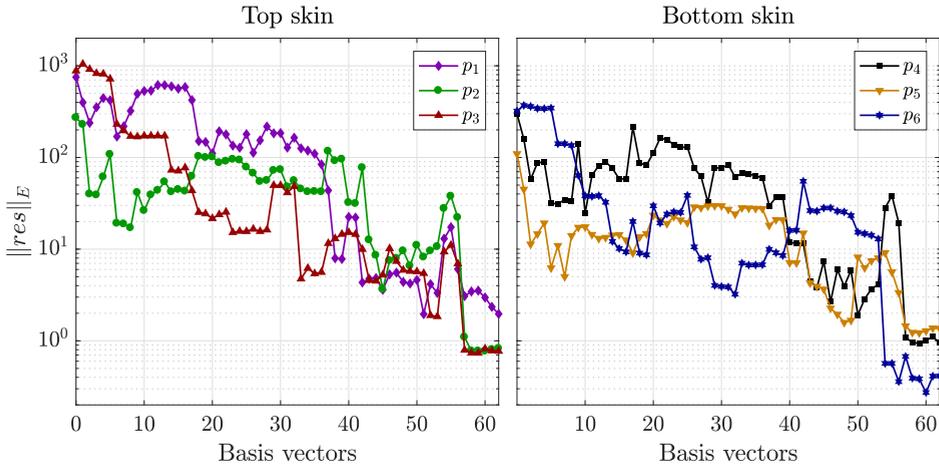


Figure 3.15: Gradient reconstruction for subcase 1 at Mach 0.72

maximum residual reduction over the entire span of basis vectors is somewhat less than 3 orders of magnitude for design region 3 and the smallest reduction is less than 1 order of magnitude for design regions 2 and 5. However, it should be mentioned that the design regions 2 and 5 start with relatively small residuals to begin with.

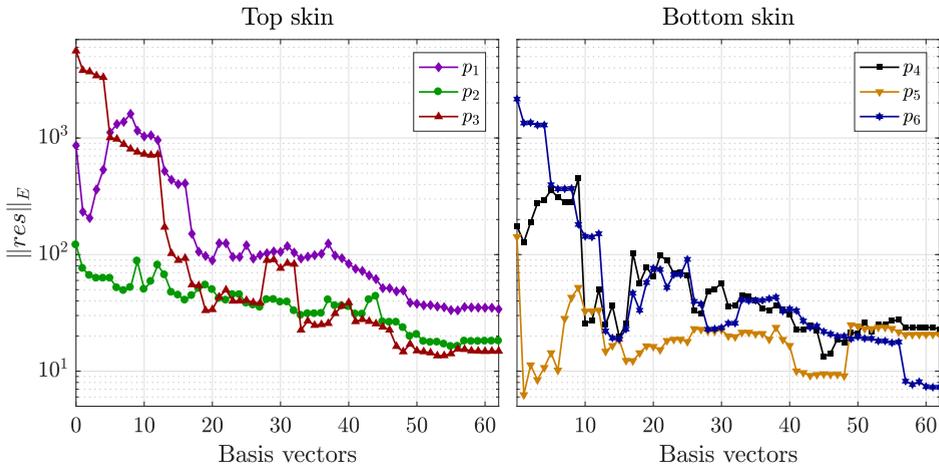


Figure 3.16: Gradient reconstruction for subcase 2 at Mach 0.81

The residuals of subcase 3, where a strong shock is present, have somewhat similar features to the previous subcase. There is a very strong initial reduction for the first 15 basis vectors as indicated by Figure 3.17. Only here, the initial residuals are much higher. The highest initial residual is 10^5 for design region 3 and the

smallest is design region 5 with 3×10^2 . The high initial residuals indicate that the VLM gradients are significantly degraded as the Mach number is increased to emulate transonic flow. Moreover, the application of the VLM in the gradient reconstruction method appears to be somewhat impeded by the presence of the recompression shock. Subcases 2 and 3 registered higher residual values with depleted basis vectors compared to subcase 1.

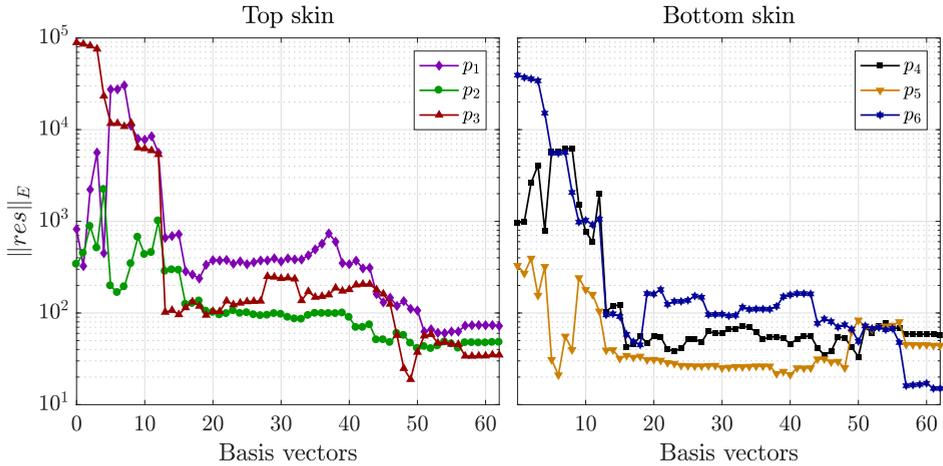


Figure 3.17: Gradient reconstruction for subcase 3 at Mach 0.87

3.5 CONCLUSIONS

A method was presented for approximating high-fidelity gradients of structural responses with respect to structural sizing variables. The method employs a vortex lattice aerodynamic model in conjunction with SVD and pseudo-inverting to reconstruct high-fidelity aero-structural Euler gradients in the transonic regime. The two case studies, of the ONERA M6 wing and the forward-swept wing, indicate that the method is capable of reconstructing the gradients as the force basis is increased. It was found that the residual between reconstructed gradients and exact (reference) gradients is not reducing monotonically. The application of the method was also extended to trimmed static sensitivity analysis for the forward-swept wing. The optimum number of basis vectors trading off computational efficiency and accuracy is case-dependent and needs to be determined once before initiating the optimization.

4

STATIC AEROELASTIC TAILORING OF A FORWARD-SWEPT WING

Forward-swept wing (FSW) aircraft have been around for a considerable time. Early, and most notable, prototypes include the Russian Belyayev DB-LK, the American Cornelius XFG-1, and the iconic German Junkers Ju 287, all developed in the early half of the 1940s. The benefits of the FSW were recognized at an early stage. The airflow over a FSW tends to migrate inward towards the root rather than outward towards the tip. This increases the prospect of attached flow at the outboard section of the wing and promotes aileron controllability at high incidence angles. Perhaps more importantly, FSWs tend to generate wash-in upon bending which further augments *aileron effectiveness*. As a result, FSW aircraft are generally more maneuverable, as they can cope with much higher incidence angles at high velocities, compared to conventional backward-swept wing configurations. Moreover, research in recent publications by Airbus [63] and the German Aerospace Center (DLR) [64] suggest that wings with low or negative sweep angles can delay the laminar-turbulent transition resulting in substantial drag reductions.

The main shortcoming of a negative sweep angle is the increasing susceptibility to structural failure by *divergence*. As the wing deflects during flight, the angle of attack increases spanwise towards the tip. This inflates the total lift on the wing, which in turn increases the tip deflection and the angle of attack until the external aerodynamic loads overcome the opposing internal forces of the structural wing-

box and the wing breaks. Early attempts to address this detrimental aeroelastic phenomenon was to enhance the rigidity of the wing, but this also resulted in considerable weight penalties which rendered any FSW design unfeasible.

In the early 1980s, with the reinvigorated surge of composite materials, efforts were made to leverage the high strength-to-weight ratio of composites to improve the structural efficiency of FSW aircraft. Krone [65] was one of the first to explore *aeroelastic tailoring* as an application of increasing the divergence speed of a FSW, while maintaining a low structural weight. His research laid the foundation for the X-29 program, which culminated in two experimental FSW airplanes built by Northrop Grumman. Designated the Grumman X-29, it started flying in 1984 and was the first FSW aircraft to break the sound barrier. Aeroelastic tailoring was a fundamental component in the design process as it enabled the divergence speed to surpass the speed of sound.

Weisshaar [66, 67] advanced the research on FSW configurations and investigated the divergence problem, bending-torsion coupling, the spanwise pressure center and aileron effectiveness using aerodynamic strip theory to predict the loads and a laminated box beam to model the structure. Librescu et al. [68, 69, 70] investigated aeroelastic tailoring using thin-walled anisotropic composite beams, and more recently Dillinger [34] investigated the effects of aeroelastic constraints such as divergence velocity, aileron effectiveness and wing twist on static aeroelastic tailoring of a FSW configuration.

The main purpose of this chapter is to monitor the effects of employing reconstructed high-fidelity aero-structural gradients in static aeroelastic tailoring of a FSW configuration. The two pertinent questions are: (i) can reconstructed gradients drive the design to an improved optimum and (ii) will the design converge faster? In an attempt to answer these questions, a design study is performed with the aeroelastic conceptual design tool PROTEUS. Load and gradient corrections are facilitated by the CFD solver ELSA. The wing configuration is the FSW introduced in Subsection 3.4.

The remainder of this chapter is organized as follows. The structural features of the wingbox, including detailed data on the mass model, are outlined in Section 4.1. Section 4.2 covers the load cases in Subsection 4.2.1, the design variables, the constraints and the objective function in Subsection 4.2.2, and the optimizer and the starting design in Subsection 4.2.3. Moreover, three optimization cases are defined in Subsection 4.2.4 to study the effects of the aerodynamic load and gradient corrections. In Section 4.3, the optimization results are analyzed and differences between the optimization cases are considered and reflected upon. A brief conclusion of this chapter is given in Section 4.4.

4.1 MODEL DESCRIPTION

The outer shape features of the FSW were described in Section 3.4 and are not repeated here. In order to reconstruct a somewhat realistic structural configuration for the purpose of static aeroelastic tailoring, a reference structural model was required. To this end, disclosed material and mass data from the FSW structural model in Dillinger [34] were used, complemented by material strength data from Werter [35]. The FSW configuration in Dillinger [34], which in turn was inspired by the DLR project *LamAiR* [71], differs only marginally in geometry compared to the FSW in the previous chapter. This allowed for a representable optimization case to be constructed.

The beam reference axis is located at 38% chord with a clamped boundary condition at the wing root. The wing is not subject to twist nor dihedral, i.e. the z-coordinates of the leading and trailing edge are consistent along the span. The aircraft mass, excluding the wing structural and non-structural masses, totals 49,500 kg. The weight of the various structural components are listed below in Table 4.1

Table 4.1: FSW aircraft mass excluding the wing (Source: Dillinger [34])

	Mass
Fuselage, Empennage, Engines	31,550 kg
Passengers	14,250 kg
Artificial Center Wing Tank (ACWT)	3700 kg
Total	49,500 kg

Apart from structural masses inherent in the finite element model, non-structural masses are accounted for as well. These comprise of structural components that are not accounted for by the wingbox finite element model and remain constant throughout the optimization, i.e. leading and trailing edge structures such as flaps, ailerons, slats, spoilers and associated actuators. These are approximated as point masses distributed in the spanwise direction along the leading and trailing edge, according to Figure 4.1. The Leading Edge (LE) and Trailing Edge (TE) mass distributions are available in Figures 4.2 and 4.3, respectively. The total sum of the LE and TE non-structural masses amount to 374.8 kg and 421.5 kg, respectively.

The wing is subject to 25 equally spaced ribs, including the ones at the root and the tip. They are aligned parallel to the free-stream and consequently to the global x-axis. The rib mass distribution is given in Figure 4.4. The top and bottom skins are reinforced in the spanwise direction by stringers. The stringers are isotropic and are subject to a constant 3% pitch with respect to the root chord. The cross-sectional area is $5 \cdot 10^{-4} \text{ m}^2$ and the longitudinal stiffness is 68.95 GPa. The density is set to 2795 kg/m^3 .

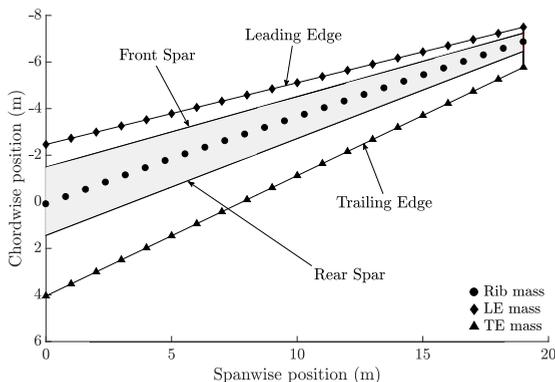


Figure 4.1: Planform view of non-structural mass locations

The ply material properties, or more specifically the elastic moduli are scaled by a factor of 0.65 in the current work. The main reason for this is to increase the flexibility of the wing and enhance the aeroelastic effects during the optimization. In contrast to the failure envelope that is constructed using strain (i) tension, (ii) compression and (iii) shear allowables for the ply, PROTEUS implements a different type of failure model based on the Tsai-Wu failure criterion. To this end, the material strength of the uni-directional AS4/3501-6 were directly extracted from Werter [35]. The material properties are available in Table 4.2.

Table 4.2: Material properties

(Source: Dillinger [34])	
Longitudinal Modulus (E_{11})	53.95 GPa
Transverse Modulus (E_{22})	5.525 GPa
Shear Modulus (G_{12})	2.73 GPa
Poissons Ratio (ν_{12})	0.35
Density (ρ)	1425 kg/m ³
UD Carbon/Epoxy (AS4/3501-6) (Source: Werter [35])	
Longitudinal Tensile Strength (X_t)	948.5 MPa*
Longitudinal Compressive Strength (X_c)	717.6 MPa*
Transverse Tensile Strength (Y_t)	23.7 MPa*
Transverse Compressive Strength (Y_c)	94.8 MPa*
Shear Strength (S)	31.6 MPa*

* Including knockdown factors for environmental effects (0.8), barely visible impact damage (0.65), and material scatter effects (0.8).

It must be mentioned that the material properties chosen from Dillinger [34] with the applied scaling factor are not in compliance with the material strength of AS4/3501-6 which is used to construct the failure criterion. This, however, was considered to be of less importance in the current conceptual design study.

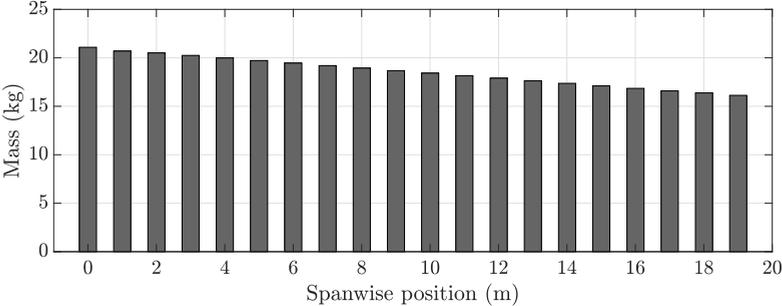


Figure 4.2: LE mass distribution (Source: Dillinger 34)

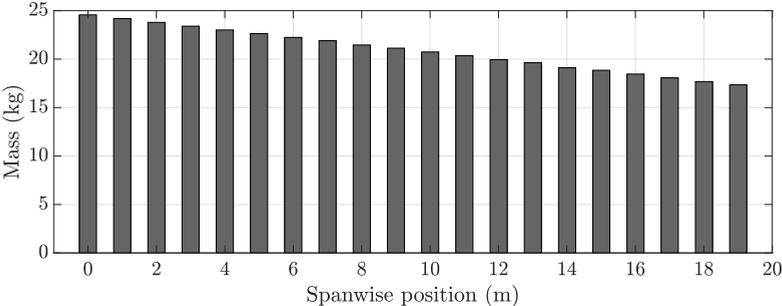


Figure 4.3: TE mass distribution (Source: Dillinger 34)

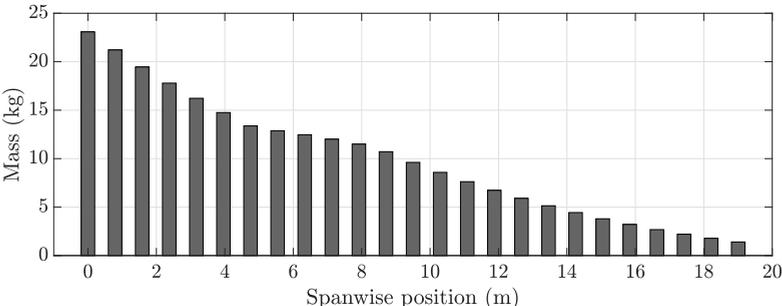


Figure 4.4: Rib mass distribution

4.2 OPTIMIZATION SETUP

In this section, the most important aspects and parameters of the optimization are described. Load cases are described in Subsection 4.2.1, while in Subsection 4.2.2 the objective function, the constraints and the design variables are outlined. The optimizer and the initial design are discussed in Subsection 4.2.3. In Subsection 4.2.4, three optimization cases are defined to investigate the influence of load corrections and gradient corrections.

4.2.1 LOAD CASES

The sizing of the wingbox, or the load carrying structure, is a process that must account for critical load cases that an aircraft might encounter during its time in service. The load cases in this thesis are limited to static manoeuvres. Dynamic manoeuvres, ground manoeuvres (taxiing or landing) and gust/turbulence encounters are consequently not considered in the optimization. The main goal in this chapter is to distinguish and examine features of the optimized results that can be attributed to higher order aerodynamics, such as recompression shocks for instance. It would therefore be superfluous to include subsonic load cases, where the aerodynamic loads of an Euler flow model are very similar to those of a VLM model. This was demonstrated on the FSW in Subsection 3.4.2 by varying the Mach number while keeping the dynamic pressure constant. It was also acknowledged in an extensive analysis by Dillinger [34] that the differences in pressure distribution between an Euler flow model and a DLM model for subsonic load cases are very small. Admittedly, subsonic load cases, if active during the optimization, will certainly affect the sizing in some regard. However, they will not contribute to the desired goal of isolating the effects between two flow models with notably different loads and gradients.

The FSW described in Sections 3.4 and 4.1 shares similar features to the FSW in the doctoral thesis by Dillinger [34]. As a consequence, the load cases in said thesis are reused in this optimization study. Subsonic sizing load cases, however, are omitted due to the aforementioned premise. In addition, a cruise load case is also omitted as it does not contribute to the structural sizing. This load case is generally included to enforce a predefined wing twist distribution at cruise condition such that maximum aerodynamic performance can be achieved. The load cases used for the current optimization study are listed in Table 4.3.

Table 4.3: Load cases

ID	Description	EAS	Altitude	Mach	n_z
1	Symm. pull up	192.5 m/s	6700 m	0.87	2.5
2	Symm. push down	192.5 m/s	6700 m	0.87	-1.0

Both load cases are subject to transonic Mach numbers. Moreover, the computation time in the current study is substantially reduced by limiting the optimization to two load cases, especially during the gradient reconstruction which is the most arduous task computationally. Introducing additional load cases would only become computationally viable if the load cases were parallelized, a feature the current PROTEUS version was not equipped with. The dynamic pressure for both load cases is equivalent to 22.7 kPa at the specified altitude and velocity. It should be mentioned that all analyses in the optimization are trimmed with gravitational effects included. This implies that the lift generated by the aerodynamic model equals half the aircraft weight multiplied by the load factor n_z .

4.2.2 OBJECTIVE, DESIGN VARIABLES AND CONSTRAINTS

The objective function of the optimization is weight minimization. The top and bottom skins of the wingbox, as well as the front and rear spars, are divided in 72 equally spaced design regions (see Figure 4.5). The top and bottom skins comprise of 24 design regions each, with 12 design regions in the spanwise direction and two in the chordwise direction. The front and rear spar comprise of 12 design regions each. The stiffness properties of the ribs and stringers are kept constant throughout the optimization and are thus not included as design variables. Each design region is subject to nine design variables: eight lamination parameters describing the in-plane and out-of-plane characteristics of composite laminates and one thickness design variable. This results in $72 \times 9 = 648$ design variables.

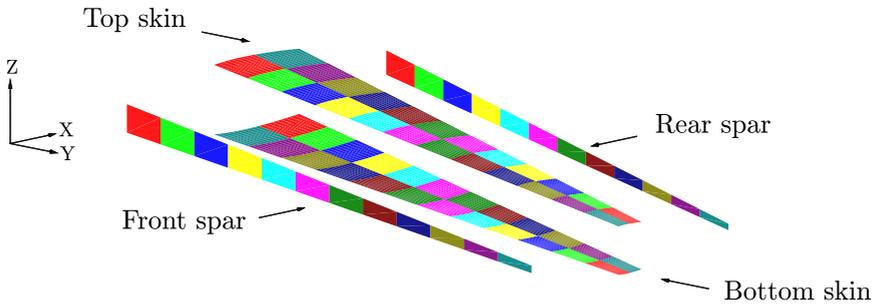


Figure 4.5: Exploded view of the design regions on the structural wingbox

The laminate thickness variation is limited to an upper and lower bound of 50 mm and 1 mm, respectively. The lamination parameters, on the other hand, are subject to a six equation feasibility constraint per design region, which totals $72 \times 6 = 432$ lamination parameter feasibility constraints over the entire design space.

Strain and buckling constraints are standard features in PROTEUS. A detailed description on these two types of structural constraints and their implementation

Table 4.4: Number of optimization constraints per load case.

Type	Number
Lamination parameter feasibility	432
Tsai-Wu strain factor	576 per loadcase
Buckling factor	5184 per loadcase
Total	432 + 5760 per loadcase

in PROTEUS can be found in Werter [35]. Here, it suffices to mention that the strain criterion is based on a Tsai-Wu strain factor where an index value above one results in structural failure, whereas the buckling constraint is based on a buckling model derived in Dillinger [34]. A buckling coefficient lower than one indicates a buckling-free design. The total number of optimization constraints per load case is listed in Table 4.4

In addition to the aforementioned constraints, PROTEUS is capable of modeling other types of constraints such as aileron effectiveness, 1G wing twist, and dynamic (flutter) and static (divergence) aeroelastic instabilities. These types of constraints, however, were inactive in the current optimization study. The gradient reconstruction scheme was only successfully implemented and verified for the strain and buckling constraints. The remaining types of constraints, which are either aerodynamic or aeroelastic, were inactive as this would require additional and extensive code modifications to the PROTEUS code.

4

4.2.3 OPTIMIZER AND INITIAL DESIGN

The Globally Convergent Method of Moving Asymptotes (GCMMA) developed by Svanberg [72] was used as a gradient-based optimizer. Since the number of constraints significantly exceeds that of the design variables, the direct method for computing gradients was applied. A feature of the GCMMA worth highlighting is that it consists of "inner" and "outer" iterations. During the latter, both the function values and associated gradients are computed. During the former, on the other hand, only the function values are evaluated. In the remainder of the thesis, if not stated otherwise, the term "optimization iteration" will refer to the outer iteration.

The initial thicknesses of the top and bottom skins, including the spars, vary linearly in the spanwise direction from 15 mm at the root to 7 mm at the tip, as is depicted in Figure 4.6. A quasi-isotropic symmetric layup $[0_{25\%}/\pm 45_{50\%}/90_{25\%}]_s$ is applied to the skins and spars, and is visualized using stiffness rosettes [36] in Figure 4.7. Moreover, the initial design is feasible, i.e. no constraints are violated.

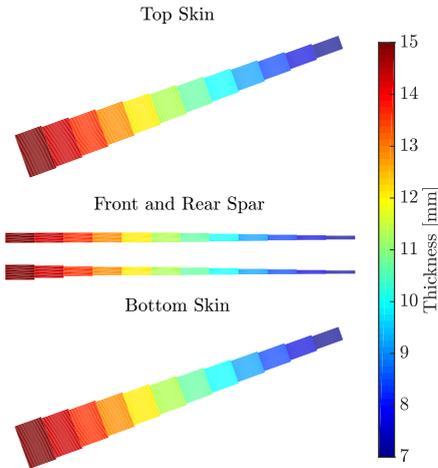


Figure 4.6: Initial thicknesses

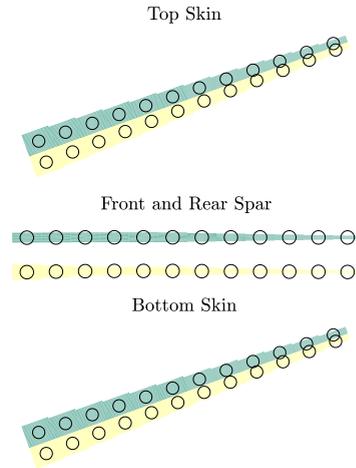


Figure 4.7: Initial stiffnesses

4.2.4 OPTIMIZATION CASES

The main interest in this chapter is to implement not only aerodynamic load corrections in static aeroelastic tailoring, but also corrected sensitivities via the gradient reconstruction scheme proposed in Chapter 3. In order to successfully study and differentiate the effects of applying aerodynamic load corrections and gradient corrections by a higher order CFD solver, three optimization cases are formulated and outlined as follows:

1. The aerodynamic loads, including the gradients, are computed by the built-in aerodynamic model in PROTEUS based on vortex lattice aerodynamics. This case is consequently denoted *VLM/VLM*. It does not require any code modification to PROTEUS.
2. The VLM aerodynamic loads are continuously corrected throughout the optimization by a higher order Euler flow model, whereas the gradients are computed by the VLM model at the static aeroelastic equilibrium determined by the Euler flow model. This case is denoted *Euler/VLM*. The methodology is similar to the CFD correction strategy presented in Dillinger [34]. The main difference is that the aerodynamic loads in the current study are corrected for each function evaluation. In Dillinger [34], the CFD corrections are only renewed every five to ten iterations and are kept constant in between iterations. The code modification effort to PROTEUS is moderate, as only load corrections are involved. The principal features of the load correction module were already outlined in Section 2.2 and in particular by Figure 2.5.

- The third case implements the gradient reconstruction strategy outlined in Chapter 3. Here, both the aerodynamic loads and the corresponding gradients are corrected by the Euler flow model. This case is consequently denoted *Euler/Euler*. Three subsets of 15, 20 and 25 basis vectors are considered to reconstruct the Euler aero-structural gradients. The three different subsets allow for a comparison of the efficiency when the vector basis is enriched. Presumably, the enrichment of the vector basis should improve the convergence rate of the optimization at the expense of an increased computational cost per optimization iteration. It demands an additional set of code modifications to PROTEUS compared to the second case in order to account for the corrected gradients.

4.3 OPTIMIZATION RESULTS

A first remark on the *VLM/VLM* optimized design can be made by observing the thickness distribution in Figure 4.8. It becomes clear that the design regions near the front spar, with the exception of the first two regions from the root, were subject to higher thicknesses. This common feature in aeroelastic tailoring shifts the elastic axis forward in order to generate wash-out (i.e. a reduced incidence angle from root to tip) upon bending. In doing so, the spanwise pressure center can be shifted inboard to produce a lower root bending moment with reduced outboard loads which results in a weight reduction. The wash-out can be visualized in Figure 4.10 for Load Case (LC)1.

4

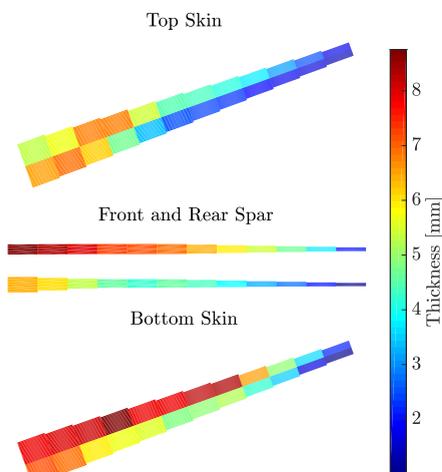


Figure 4.8: Converged thicknesses

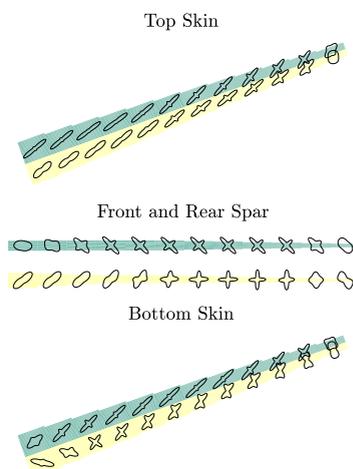


Figure 4.9: Converged stiffnesses

A second feature which stimulates weight minimization, by alleviating the root

bending moment, is the stiffness direction of the design regions. The stiffness distribution is governed by the lamination parameters and can be visualized in Figure 4.9. The regions close to the root (see top skin) are designed with a high in-plane stiffness in the spanwise direction, while the dominant in-plane stiffness close to the tip is gradually tilted towards the flow direction. This stiffness configuration promotes a negative bend-twist coupling [34, 73], i.e. as the wing bends upwards the tip will twist nose down. This results in an enhanced wash-out effect, amplifying the effect due to the forward shift of the elastic axis, moving the spanwise center of pressure inboard.

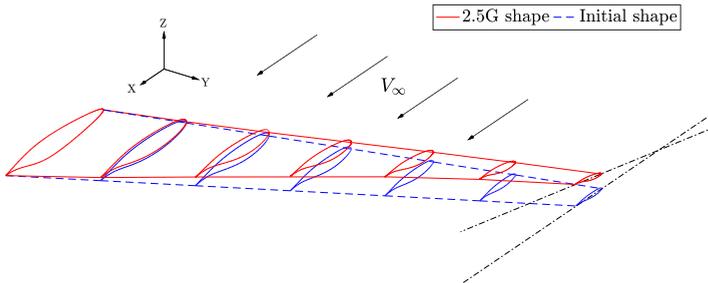


Figure 4.10: 3D visualization of the wash-out

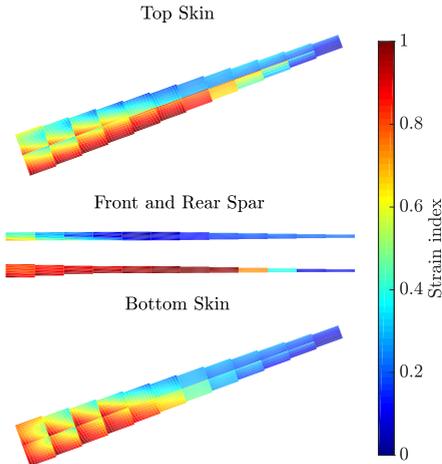


Figure 4.11: Strain LC1

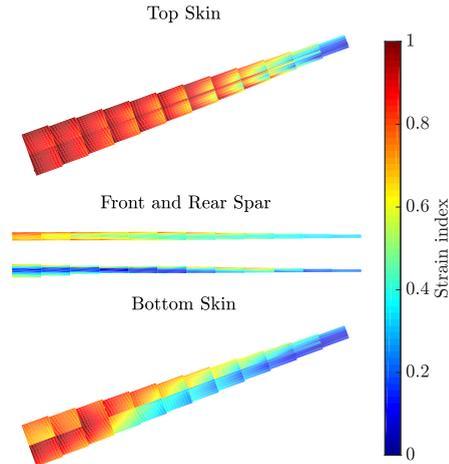


Figure 4.12: Strain LC2

All the constraints described in Subsection 4.2.2 were active during the current optimization, i.e. the failure index was equal to, or very close to, one on at least one of the 72 design regions. The strain constraint appeared to be more active for LC2 on the top skin compared to LC1 on the bottom skin (see Figures 4.11-4.12).

The buckling constraint was clearly more dominant for LC2 in the region close to

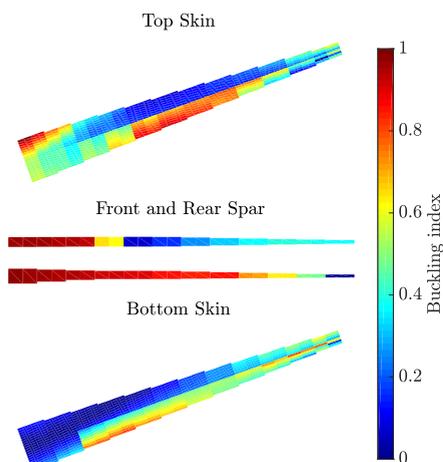


Figure 4.13: Buckling LC1

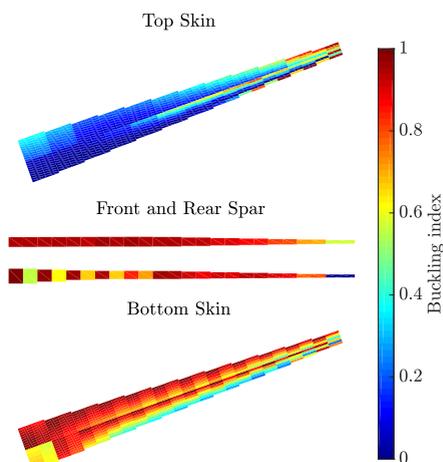


Figure 4.14: Buckling LC2

the root on the bottom skin, as can be seen in Figure 4.14. Moreover, the spars, front and rear, noticeably hit the buckling constraint limit for both load cases. However, the buckling constraint did not appear to be very active on the top skin for LC1 (see Figure 4.13).

The fact that LC2 appears to trigger more design constraints for both failure and buckling compared to LC1 indicates that the vertical tip deflection is larger for LC2. Indeed, LC1 and LC2 resulted in a 0.90 m and a -2.87 m vertical tip deflection, respectively. That is, LC2 produced more than a threefold larger absolute vertical tip deflection compared to LC1. It was anticipated that the reduced stiffnesses in Table 4.2 would trigger large deformations (in fact this was the desired outcome). However, the substantially larger deformations, and the corresponding strain and buckling responses, for LC2 compared to LC1 were unforeseen. As such, to better comprehend the characteristics of the current optimization, and in particular the influence of the airfoil shape, additional optimizations were necessary. This is discussed in detail in the following subsection.

4.3.1 AIRFOIL INFLUENCE

To investigate the impact of the airfoil shape on the optimized results, two additional optimization cases were performed. The first was identical to the *VLM/VLM*, only with a supercritical OAT15A airfoil (see Figure 3.9) along the entire span. This case was denoted *OAT15A* and the intention here was to shed some light on the peculiar characteristics of the airfoil shape close to the root, as is depicted in Figure 4.8. The second optimization, denoted *NACA0012*, had a symmetric NACA0012 airfoil adopted along the entire wingspan, whereas all

remaining parameters were kept identical to the *VLM/VLM* case. The intention of the second optimization was to investigate the influence of the airfoil shape on the vertical tip deflection. It should be mentioned that the NACA0012 airfoil was selected because of its generic character and absence of camber. Any other symmetric airfoil would have been equally sufficient as the vortex ring elements in PROTEUS would have been modeled in the same manner, i.e. by distributing them on the camber surface which in the case of a symmetric airfoil overlaps with the symmetric surface.

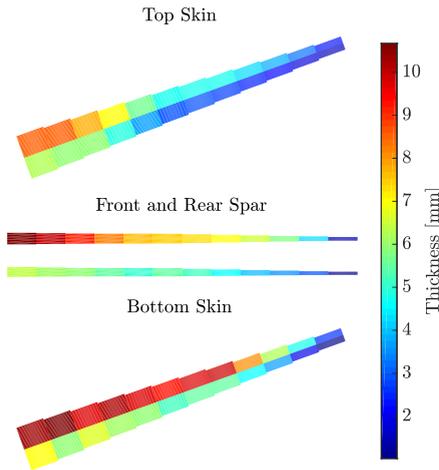


Figure 4.15: *OAT15A* thicknesses

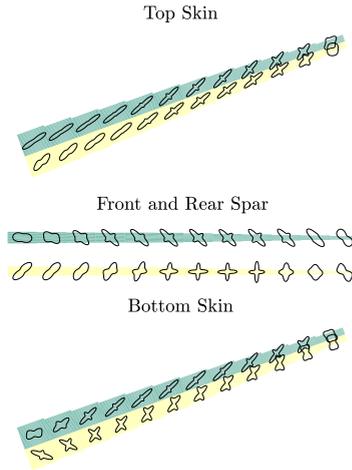


Figure 4.16: *OAT15A* stiffnesses

In the first case, it became evident that the airfoil shape has a substantial impact on the optimization and the converged thickness distribution, as can be seen by comparing Figures 4.8 and 4.15. By applying a *OAT15A* airfoil along the entire span, the design region thicknesses at the wing root are increased more towards the front spar. This contradicts the previous results in Section 4.3, where the *OALE10* airfoil at the root triggered a shift of increasing thickness towards the rear spar. It can therefore be derived that the main source of this spanwise shift in thickness distribution can be attributed to the change in the airfoil shape. In contrast to the thickness distribution, there does not seem to be an equally significant change in the stiffnesses when comparing Figures 4.9 and 4.16. The strain and buckling responses for the *OAT15A* optimization were not selected for presentation due to the similarities with the *VLM/VLM*.

In the second case, the *NACA0012* optimization indicated a similar thickness distribution pattern to that of the *OAT15A*, i.e. a increasing thickness at the front spar for the entire wing span (see Figure 4.17). What is of greater interest here is that LC1 and LC2 converged to a 2.60 m and a -1.05 m vertical tip deflection, respectively. This opposes the previous results in Section 4.3, where LC2 resul-

ted in a significantly higher absolute tip deflection compared to LC1. Moreover, an additional +1G load case was included in the *NACA0012* optimization which resulted in an identical absolute tip deflection as LC2. This validated the optimization because the response has to be symmetric in case of a symmetric airfoil. However, the +1G load case did not affect the structural sizing process since the strain and buckling responses generated by LC1 supersede those generated by the +1G load case. Based on the results, it can be concluded that the particular airfoil shape was the primary source for the large negative tip deflection of LC2 in the *VLM/VLM* case, which in turn triggered the response illustrated in Figures 4.12 and 4.14. An additional feature of influence that merely amplified the difference in tip deflection between LC1 and LC2 is the reduced stiffness values in Table 4.2, which were judiciously selected to augment the aeroelastic response.

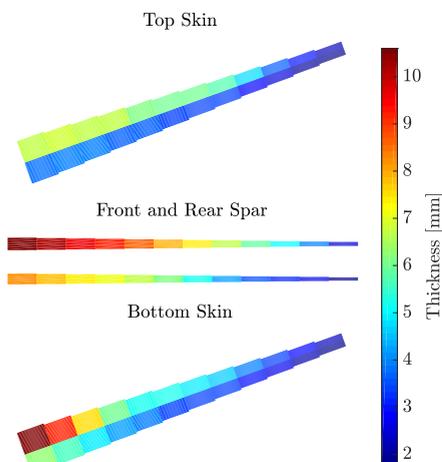


Figure 4.17: *NACA0012* thicknesses

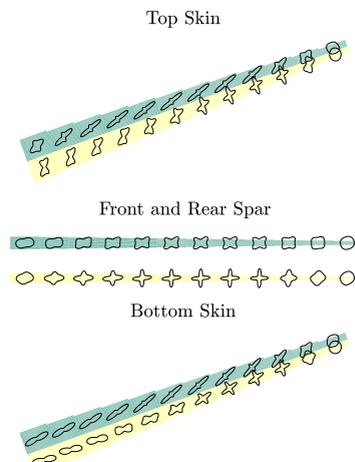


Figure 4.18: *NACA0012* stiffnesses

Compared to the *VLM/VLM* case, and in agreement with the converged vertical tip deflections, the strain and buckling response for the *NACA0012* optimization were most prominent for LC1 on the bottom skin and top skin, respectively. The results are illustrated in Figures 4.19, 4.22.

Before concluding the current subsection, a comparison of the spanwise lift and moment distributions was made. Figures 4.23 and 4.24 depict the initial and optimized trimmed aeroelastic solutions for LC1 and LC2, respectively. The lift was computed by summing up the chordwise contributions on each of the 47 spanwise panel sections, hence the 47 markers in the aforementioned figures. The aerodynamic moment was computed about the local quarter chord of each spanwise panel section.

It became clear that the *OAT15A* case generated more lift close to the root compared to the *VLM/VLM*, but that the chordwise center of pressure was located

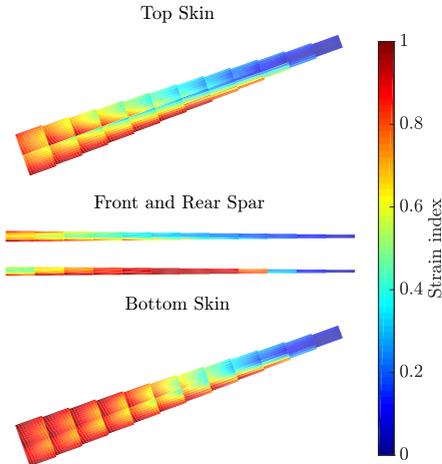


Figure 4.19: Strain LC1 (*NACA0012*)

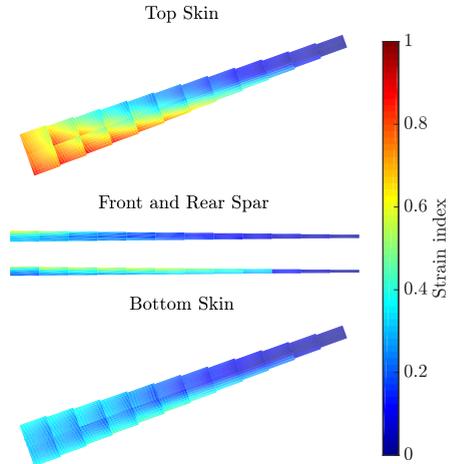


Figure 4.20: Strain LC2 (*NACA0012*)

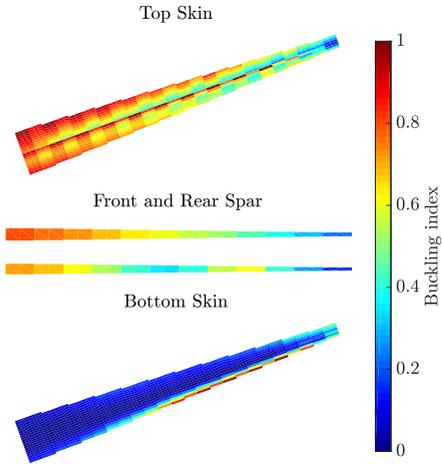


Figure 4.21: Buckling LC1 (*NACA0012*)

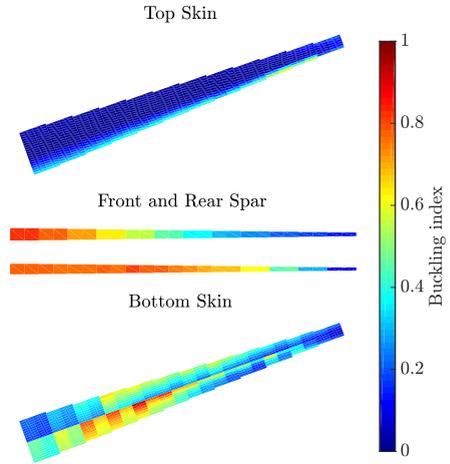


Figure 4.22: Buckling LC2 (*NACA0012*)

4. STATIC AEROELASTIC TAILORING OF A FORWARD-SWEPT WING

closer to the trailing edge, subsequently prompting a much larger negative torsion moment. This is the reason why the thickness distribution is different between *VLM/VLM* and *OAT15A* in the root region. Another interesting note to make is that the spanwise pressure center, distinguished in the figures by a filled marker, was shifted closer to the root for all three cases. This can be directly associated to the objective function which is to minimize the weight by reducing the root bending moment. It is quite remarkable that the reduction in root bending moment not only triggered a shift of the spanwise pressure center towards the root by increasing the lift in the said region, but also by generating negative lift at the tip. This can be seen in Figure 4.23 for the *VLM/VLM* and *OAT15A*. Notice how the lift decreases almost linearly from the root to the tip, whereas the optimized designs resemble a reversed S-shaped lift distribution.

Regarding the *NACA0012* case, a nearly zero quarter chord pitching moment can be observed for both initial as well as optimized designs. This was an expected outcome since an uncambered flat plate does not generate pitching moment around the quarter chord, even when the angle of attack and lift vary [74]. Moreover, although the initial spanwise pressure center for LC2 is at 15 m, the corresponding optimized pressure center is located at 7 m. This can be compared with the two remaining cases where the pressure center of the initial designs are at roughly 14–15 m and the optimized designs are at 12–13 m. This observation underscores why the negative tip deflection of the *NACA0012* case was not as substantial as the remaining two cases.

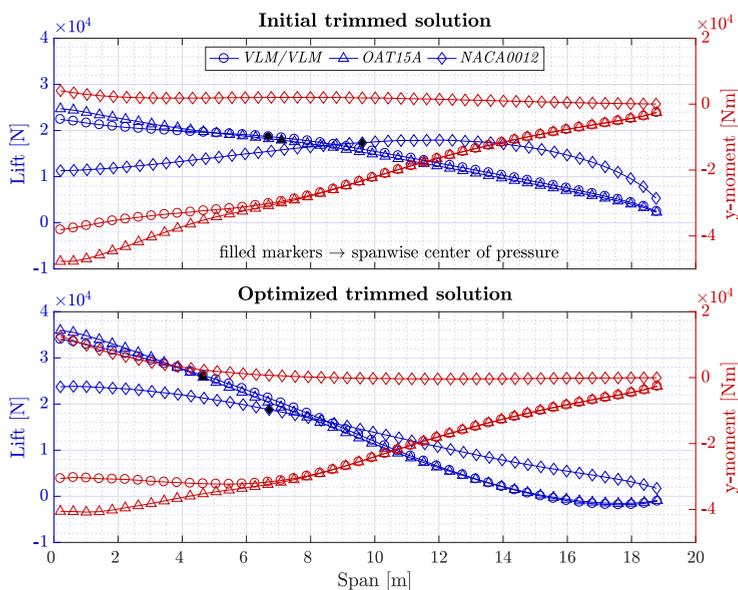


Figure 4.23: Spanwise lift and y-moment for LC1 at initial and optimized trimmed solution

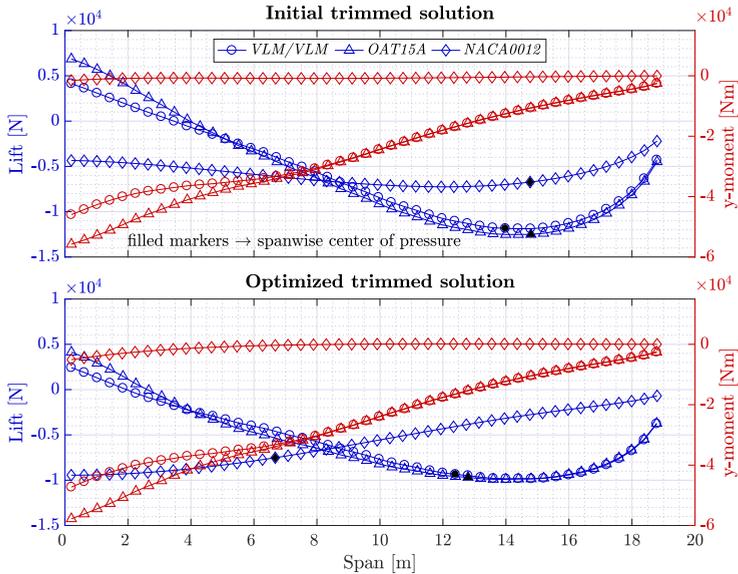


Figure 4.24: Spanwise lift and y-moment for LC2 at initial and optimized trimmed solution

4.3.2 EULER LOAD AND GRADIENT CORRECTIONS

In Section 3.4, the gradient reconstruction method was evaluated for Mach = 0.72, 0.81 and 0.87 at fixed dynamic pressure. Indeed, it was shown that the first subcase of Mach = 0.72 did not produce a recompression shock, whereas Mach = 0.87 generated a strong shock near the trailing edge (see Figure 3.13). To verify this, and to better comprehend the limitations of linearized potential flow methods, a simple case was constructed. The FSW was analyzed at its initial configuration (i.e. without any deformation to the structure) at five degrees angle of attack. The dynamic pressure was kept constant at 22.7 kPa (the same as the optimization) and a steady computation was performed by the VLM and the Euler flow solver at Mach = 0.60 and 0.87. Here, we chose a slightly lower Mach number for the subsonic case, merely to increase the contrast between subsonic and transonic flow. The cross-sectional pressure difference at various spanwise stations can be observed in Figures 4.25 and 4.26.

The ΔC_p at Mach = 0.60 indicated a good agreement between the two aerodynamic methods. There was a small offset on the VLM results as it appeared to underpredict the aerodynamic pressure. This is in agreement with the results obtained by Dillinger 34 and can be attributed to the thin-plate approximation which is employed by PROTEUS (i.e. the vortex elements are distributed on the camber surface and not on the actual wing surface). Moreover, the VLM results lacked in resolution at the LE and TE and could be improved by increasing the

4. STATIC AEROELASTIC TAILORING OF A FORWARD-SWEPT WING

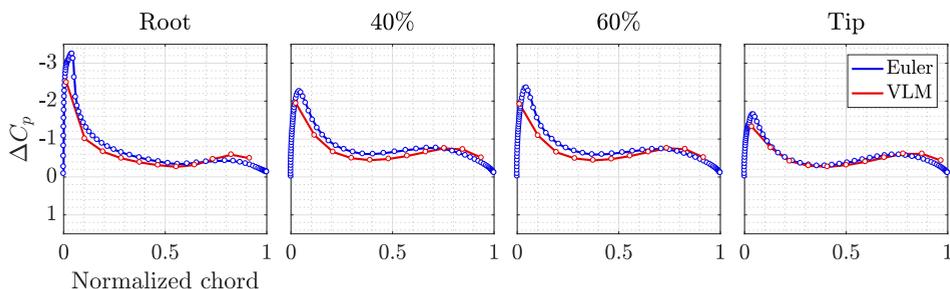


Figure 4.25: Cross-sectional ΔC_p at Mach = 0.60 and five degrees incidence angle

chordwise panel density. Unfortunately, increasing the number of panels resulted in prohibitively high computation time, in particular when the vortex lattice aerodynamic stiffness matrix $\tilde{\mathbf{K}}_a$ was assembled. Hence, the number of panels in the chordwise direction was limited to eleven in this study.

4

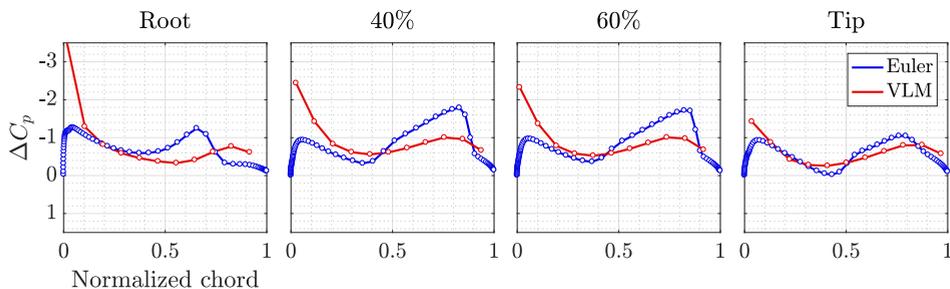


Figure 4.26: Cross-sectional ΔC_p at Mach = 0.87 and five degrees incidence angle

The aerodynamic methods display a different behavior in transonic flow, as can be seen in Figure 4.26. The pressure difference of the Euler flow changed significantly and a strong shock was manifested in the form of a pressure peak near the trailing edge. The VLM results, on the contrary, were much smoother and resembled the VLM results from the subsonic case, save the high pressure peak at the leading edge. The said pressure peak was not due to a recompression shock, as it was already highlighted that linearized potential flow methods lack this predictive capability, but due to the Prandtl-Glauert correction mentioned briefly in Subsection 2.1.1. The Prandtl-Glauert correction is a method used to account for compressibility effects in high-subsonic flows. It is based on a simple coordinate transformation featured by a coefficient η [37]:

$$\eta = \sqrt{1 - M_\infty^2} \quad (4.1)$$

where M_∞ is the free-stream Mach number. The lift coefficient, for instance,

corrected by the Prandtl-Glauert rule, then becomes

$$C_L^{M>0} = \frac{C_L^{M=0}}{\eta} \quad (4.2)$$

where $C_L^{M>0}$ and $C_L^{M=0}$ are the lift coefficients with and without correction, respectively. It can be derived from Equations 4.1 and 4.2 that as the flow approaches the speed of sound (i.e. $M_\infty \rightarrow 1$), the corrected lift coefficient goes to infinity (i.e. $C_L^{M>0} \rightarrow \infty$). As a consequence, the Prandtl-Glauert correction tends to overpredict aerodynamic loads at transonic flow [37]. A Mach dependency study, where the lift coefficient is plotted as a function of Mach number, can typically shed some light on the range in which the Prandtl-Glauert correction can be considered valid for a certain configuration. Such a study was not conducted in the current optimization since the Mach number for the two load cases in Table 4.3 was well beyond the limit of linearized potential flow methods.

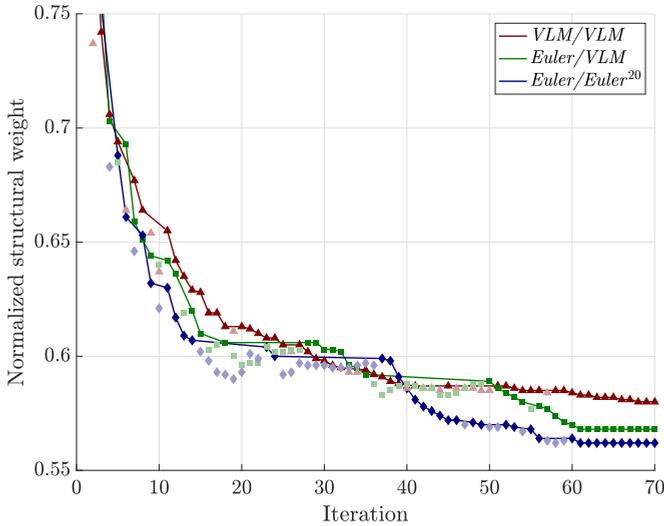


Figure 4.27: Convergence history of optimization cases

The convergence history of the three optimization cases, outlined in Subsection 4.2.4, is illustrated in Figure 4.27. For the purpose of increasing the readability, the y-axis was subject to an upper and lower limit of 0.75 and 0.55, respectively. The *VLM/VLM* case noticeably resulted in the smoothest convergence. Apart from iterations 3, 7, 11 and 20, the structural weight appeared to reduce monotonically until convergence. The optimization was not subject to a convergence criterion in this study. Rather, the optimization was allowed to advance for a considerable amount of iterations until it could be established, through a visual representation in Figure 4.27, that convergence was achieved, whereupon the optimization was terminated.

The *Euler/VLM* optimization demonstrably converged to a lower structural weight than the *VLM/VLM*. This observation was also made by Dillinger [34], where the optimization including Euler load corrections with unbalanced laminates resulted in a lower structural weight compared to the optimization excluding Euler load corrections. Considering the fact that the analyses were trimmed and that the total lift for all cases were identical, this outcome can be attributed to the disparity in load distribution between the two aerodynamic models in transonic flow. Indeed, as Figure 4.26 indicates, the recompression shock triggered a rearwards shift of the chordwise center of pressure towards the TE, effectively increasing the negative torsional moment. This generated a wash-out effect and reduced the need for the optimizer to compensate by adding mass at the LE. The chordwise center of pressure of the VLM, on the other hand, is comparably closer to the LE, thus prompting the optimizer to add weight at the LE to generate wash-out and stimulate a reduced root bending moment. Hence, the load distribution of the VLM increases the weight penalty associated to load alleviation in transonic flow for this particular case.

The *Euler/Euler*²⁰ case, where the superscript signifies the number of basis vectors employed during the gradient reconstruction, displayed an improved convergence for the initial 19 iterations. The convergence then stagnated for the subsequent 19 iterations, without any notable reduction in the structural weight, followed by a rapid decrease in the objective function after iteration 38. The optimization eventually converged to a lower structural weight than the *Euler/VLM*. It is interesting to note that the introduction of corrected gradients prompted a distinctly different convergence path. However, it did not contribute to a monotonic decrease in structural weight, nor did it converge to a substantially different optimum. A possible, and conceivable, source for this outcome is that the reconstructed gradients were not sufficiently accurate or consistent with the structural response. To this end, as was declared previously, two additional optimizations with different sets of basis vectors were performed to investigate the difference in convergence behavior.

In Figure 4.28, a convergence study was performed for different sizes in the vector basis. It can be observed that the case with the smallest subset of 15 basis vectors, *Euler/Euler*¹⁵, encountered substantial convergence difficulties after the first 20 iterations. It did not fully converge within 70 iterations and even resulted in a design with a higher structural weight than all the preceding cases that were depicted in Figure 4.27. This implies that the vector basis was not sufficiently large to reconstruct accurate Euler aero-structural gradients. Conversely, the case with the largest subset of 25 basis vectors, *Euler/Euler*²⁵, displayed an improved convergence behavior after the first 20 iterations. Although it converged to the same structural weight as *Euler/Euler*²⁰, the convergence appeared to be more rapid in the range between iterations 24 and 38. It should be mentioned that 25 basis vectors was an upper limit in this thesis. Optimizations with exceeding subsets resulted in excessive computation costs and were consequently not considered.

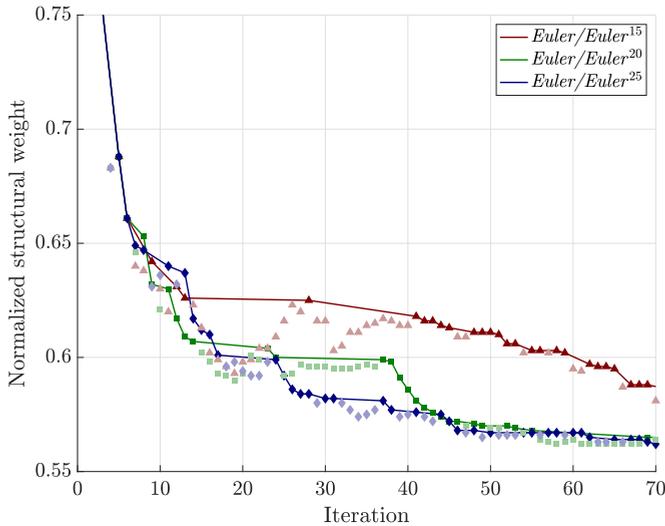


Figure 4.28: Convergence history of optimization cases with varying number of basis vectors

4.4 CONCLUSIONS

In this chapter, a structural sizing study was defined to investigate the potential of the gradient reconstruction scheme developed in Chapter 3. The optimization was performed on a FSW configuration and the in-house conceptual design tool PROTEUS was employed to parametrize the structural layout of the wingbox. The CFD solver ELSA, delivered corrected loads and associated gradients by a 643K cell Euler model. Three optimization cases were defined: (i) one with loads and gradients generated by the built-in VLM in PROTEUS, (ii) one with corrected loads by the Euler flow model but with VLM gradients and (iii) one with corrected loads and gradients by the Euler flow model, employing the methodology in Section 3.2. The three cases converged to different optima, where the third case including corrected gradients converged to the lowest structural weight. The limited applicability of VLM to shock-free flow conditions emphasizes the importance of introducing high-fidelity aerodynamics in structural sizing, not only in the form of corrected aerodynamic loads but also rectified aero-structural gradients. In Section 4.3, the convergence characteristics were shown to improve for the case with Euler gradients as the subset of basis vectors was enriched. Conversely, the optimized result deteriorated when the subset was limited to a small set of vectors.

It must be stressed that the Euler equations, which are used extensively in this design study, do not constitute the highest order of aerodynamic equations and must be considered carefully in structural sizing. The deployment of even higher order CFD, such as the RANS equations, can result in considerable shock reloca-

4. STATIC AEROELASTIC TAILORING OF A FORWARD-SWEPT WING

tion, effectively altering the pressure distribution and in extension the converged design. Additionally, a convergence study on the mesh size of the Euler grid should be performed to affirm the reliability of the aerodynamic loads generated by the Euler equations. Since the mesh was delivered as an input from ONERA, this was not possible in the current study.

5

CONCLUSIONS AND RECOMMENDATIONS

Panel methods are extensively used throughout industry and academia in conceptual design. They provide rapid results, but their applicability is limited up to high subsonic flow. Modern-day commercial aircraft predominantly operate in transonic flow, and it is therefore important to include high-fidelity aerodynamic methods in the design process that provide reliable results in this flight condition.

The primary motive, at the outset of this dissertation, was to advance the research efforts of Dillinger [34] to not only account for high-fidelity aerodynamics in static aeroelastic tailoring, but to account for high-fidelity aero-structural gradients as well. The main objectives in this thesis were defined in the introduction and are repeated below for the sake of clarity:

- 1. Develop a method and establish a framework for solving direct high-fidelity aero-structural gradient problems.*
- 2. Develop a method for computing structural response gradients in high-fidelity aero-structural sensitivity analysis.*

In the following section, the chapters are reviewed and the successfulness of accomplishing the main objectives is reflected upon. In Section 5.2, recommendations for future research are discussed.

5.1 CONCLUSIONS

Prior to solving the aero-structural gradient problem, the underlying governing equations must be solved. Indeed, the initial challenge in this thesis was to solve the high-fidelity static aeroelastic problem. The problem was partitioned in computational modules to make it a more tractable undertaking. Each of the modules was assigned individual subtasks with information exchanged in a systematic fashion by a coupling module. The modules were: (i) two aerodynamic solvers (one considered medium-fidelity and a second considered high-fidelity), (ii) a structural solver, (iii) a coupling module and (iv) an aerodynamic stiffness matrix module.

The high-fidelity aerodynamic solver used throughout this thesis was ELSA. Although it is capable of solving both the Euler and RANS equations, it was primarily used in this thesis to solve the former, with an exception in chapter two where the RANS equations were implemented as well. The limited use of the RANS equations can be attributed to the substantial increase in computation cost and the unavailability of a structured RANS mesh for the forward-swept wing model in chapters three and four. The medium-fidelity aerodynamic solver was the in-house aeroelastic design tool PROTEUS, solving the linearized potential flow equations with a Prandtl-Glauert compressibility correction. The structural solver in chapter two was NASTRAN, while the structural solver in chapter three was later changed to PROTEUS in order to accommodate for the subsequent optimization in chapter four. The coupling module was designed to account for load transfer from the aerodynamic surface mesh to the structural degrees of freedom, but also to conform the aerodynamic grid to the structural deformations. To that end, a simple nearest-neighbor algorithm was adopted for the load transfer, whereas a radial basis function approach was adopted for the mesh deformation. In particular, a globally supported thin-plate spline function was used to determine the translation of the aerodynamic grid points. To reduce the computation time, the mesh deformation routine was parallelized with the CPU time shown to scale linearly up to 16 processors. The aerodynamic stiffness matrix, describing the change in aerodynamic loads on the structural grid due to perturbations on the structural degrees of freedom, was constructed by PROTEUS.

With all necessary modules available, the high-fidelity static aeroelastic equations were efficiently solved by a defect-correction methodology. This approach, previously developed and described by Dillinger [34], was successfully implemented in this thesis to converge the coupled nonlinear problem to a state of static equilibrium.

The solution of the static aeroelastic problem provided the necessary inputs for the subsequent aero-structural gradient problem. In chapter two, a new method was proposed to solve the gradient problem using the aerodynamic stiffness matrix as a preconditioner to accelerate the convergence. This method was benchmarked

against a conventional linear block Gauss-Seidel (LBGS) method with and without under-relaxation. The model which was used to study the convergence rate was the ONERA M6 wing in transonic flow condition. The proposed method was shown to outperform the LBGS by up to 41% in computation time in conditions of strong aeroelastic coupling using the Euler flow model, whereas a 16% reduction in computation time was registered for the same conditions using the RANS flow model. The proposed method demonstrated promising results for the ONERA M6 wing, but admittedly requires further evaluation on more conventional aircraft configurations such as the Common Research Model. Nevertheless, a high-fidelity aero-structural gradient computation framework was at hand and the first research objective of the thesis was considered successfully accomplished.

As outlined in the introduction, the contemporary state-of-the-art method in aero-structural gradient computation for MDO relies on constraint aggregation. Structural constraints are lumped in one or a limited number of aggregation functions, whereupon the gradients are efficiently determined by solving the coupled adjoint problem. However, this methodology delivers gradients that guides the optimizer to suboptimal designs and it was therefore deemed important to provide an alternative method of obtaining said gradients.

In chapter three, a new method was presented to reconstruct high-fidelity aero-structural gradients using the aerodynamic stiffness matrix from chapter two. The method employed singular value decomposition on the aerodynamic stiffness matrix to inexpensively construct a set of basis vectors, which were then used to solve the high-fidelity aero-structural gradient problem and subsequently reconstruct the corresponding high-fidelity aerodynamic stiffness matrix by pseudo-inverting. The method was tested in two case studies: (i) the Onera M6 wing for a fixed angle of attack and (ii) a forward-swept wing model in a trimmed condition provided by ONERA. By benchmarking the reconstructed gradients with accurate semi-analytical gradients, the method demonstrated successful high-fidelity aero-structural gradient reconstruction with an increasing subset of basis vectors. Moreover, the spectrum of the forward-swept wing case study was extended by investigating the efficiency of the proposed method for three Mach numbers: 0.72, 0.81 and 0.87. It should be mentioned that the dynamic pressure was kept fixed for all three cases. The first Mach number, corresponding to high-subsonic flow, resulted in very promising reconstruction, while the reconstruction for the remaining two Mach numbers, corresponding to transonic flow, indicated that the quality of the basis vectors provided by the VLM decays with increasing Mach number. Nevertheless, gradient reconstruction was still possible for very high Mach numbers. Unfortunately, a monotonic decrease of the error with an enrichment of the vector basis could not be guaranteed. Consequently, the method in its current formulation makes it difficult to estimate how many basis vectors would be required for a sufficient reconstruction of the high-fidelity gradients. Regardless, the second research objective was successfully accomplished by establishing a new method to compute structural response gradients in high-fidelity aero-structural

sensitivity analysis.

In chapter four, the applicability of the gradient reconstruction method was investigated for static aeroelastic tailoring. Structural weight minimization was performed on the forward-swept wing from chapter three. The optimization considered buckling and material failure constraints, but disregarded other important constraints such as divergence and aileron effectiveness. Moreover, to minimize the computation time, the optimization was limited to two transonic maneuver load cases, while the design variables were lamination parameters and laminate thicknesses. Three optimization cases were defined to study the effects of introducing high-fidelity load corrections: (i) one with loads and gradients generated by the built-in VLM in PROTEUS, (ii) one with corrected loads by an Euler flow model using ELSA but with VLM gradients delivered by PROTEUS and (iii) one with corrected loads and gradients by an Euler flow model using ELSA. The convergence behavior was studied for the three cases and comparisons were made. It was demonstrated that the second and the third case converged to a lower structural weight by virtue of corrected aerodynamic load distribution. Since the static aeroelastic analyses were trimmed for all three cases, the total lift did not differ, but only the distribution of aerodynamic pressure exerted on the structural wing-box. The third case, including corrected gradients, displayed superior convergence characteristics compared to the second case, in particular when the subset of vectors was enriched for the gradient reconstruction. However, the concession for improved convergence using the gradient reconstruction method in the third optimization case is the increased computation cost.

5.2 RECOMMENDATIONS

Several recommendations are brought forward here to inspire and further promote the topic of aero-structural gradient computation in MDO. The gradient reconstruction method proposed in chapter three appeared as a promising alternative to the conventional approach, which is combining constraint aggregation with the adjoint method. However, a concession of the proposed method in its current implementation is its inability to enforce a monotonic reduction of the error as the vector basis is enriched. This shortcoming makes it difficult to determine a minimum number of basis vectors required for a sufficient reconstruction of the aero-structural gradients. Moreover, benchmarking the proposed method against the aforementioned conventional method could also provide some insight on its true potential. A wing-body configuration, for instance the Common Research Model, with a larger set of coupling nodes would provide a good test case of the proposed methods ability to reconstruct accurate gradients with a limited number of basis vectors using a linearized potential flow model.

In this thesis, the compressible Euler equations were predominantly used as a

high-fidelity model to determine nonlinear aerodynamic flow, including recompression shocks, in the transonic regime. However, modeling recompression shocks without the consideration of viscosity can substantially diminish the credibility of the shock location and in extension the optimized results. To account for shock-induced flow separation and by that obtain a superior prediction on the shock location, the RANS equations are recommended. However, extending the optimization to include viscous effects by adding boundary layers in the computational model requires additional attention to the mesh deformation strategy. While the radial basis function approach is a robust deformation method for Euler grids, this is not generally the case for RANS grids. The RBF approach implemented in this thesis considers only linear elastic deformations where the displacements in the three coordinate axes are independent of one another. It is therefore not the most suitable mesh deformation method when local rotations are significant. For instance, a highly deformed RANS mesh can result in very skewed boundary layers which makes it difficult for the solver to reduce the flow residuals and converge to a steady-state. To address this shortcoming, local rotations must be considered by a high-quality mesh deformation strategy which preserves the orthogonality of the undeformed cells.

Considering the computational bottleneck during the optimization in chapter four, the gradient reconstruction method could be parallelized to improve the optimization time. In fact, several instances of parallelization could be accomplished. In its current layout, only the aerodynamic gradient computation in ELSA and the mesh deformation were parallelized. Additional code parallelization could be applied to the load cases, such that each load case is distributed and computed in parallel by a separate set of processors. The reduction in computation time for the optimization would then roughly be proportional to the number of load cases. Another improvement can be made by allowing the basis vectors to be computed in parallel during the gradient reconstruction. For instance, there are very efficient block-Krylov algorithms that can account for multiple right-hand sides during the iterative solution of the linear systems.

Only two types of structural constraints were applied during the optimization in chapter four: strain and buckling failure. Other constraints of static aeroelastic nature, such as divergence and control effectiveness, were not included. However, these constraints are extremely important during structural design and the certification process and must be considered in future design studies. Indeed, divergence and control effectiveness are of particular interest in forward-swept wing configurations and generally have a significant impact on the structural design. All of the aforementioned aeroelastic constraints were available in PROTEUS by virtue of the built-in unsteady and steady vortex lattice method but were inactive during the optimization due to the inconsistency with the load and gradient corrections. For instance, the structural constraints were determined by the aerodynamics of the Euler flow model, but the response and gradients of the control surfaces, if active, would have been governed by the VLM. It is therefore sugges-

5. CONCLUSIONS AND RECOMMENDATIONS

ted that the aeroelastic constraints are accounted for by the Euler flow model as well such that they become consistent with the structural responses.

Finally, only structural optimization with static aeroelastic deformations was considered in this thesis. That is, the analysis and sensitivity analysis were multi-disciplinary, but the optimization was mono-disciplinary. To achieve MDO which is an aspiration in the industrial design process, the shape of the wing must be allowed to change with an objective function that, in addition to structural weight minimization, accounts for a reduction of the aerodynamic drag. To accomplish this, a parametrization of the wing shape is required to facilitate the perturbation of the surface grid based on the design variable output from the optimizer. For instance, this can be accomplished by a non-uniform rational basis spline (NURBS) mathematical model. Moreover, an objective function is required that accounts for a simultaneous improvement of both structural and aerodynamic efficiency. The Breguet range equation is just one example of such an objective function formulation. In addition to these modifications, cross-coupling partial derivatives are required that describe how perturbations of structural variables affect aerodynamic residuals, and vice versa.

BIBLIOGRAPHY

- [1] Boeing, Current market outlook 2017-2036 (2017).
- [2] Airbus, Global market forecast 2017-2036: Growing horizons, ISBN: 978-2-9554382-2-6 (2017).
- [3] European Commission, Flightpath 2050: Europe's vision for aviation, EUR 098 EN (2011).
- [4] A. Kharina, D. Rutherford, Fuel efficiency trends for new commercial jet aircraft: 1960 to 2014, Tech. rep., International Council on Clean Transportation (ICCT), Washington DC, USA (2015).
- [5] A. Jameson, Re-engineering the design process through computation, *Journal of Aircraft* 36 (1) (1999) 36–50. [doi:10.2514/2.2412](https://doi.org/10.2514/2.2412).
- [6] W. G. Roeseler, B. Sarh, M. U. Kismarton, Composite structures: The first 100 years, in: 16th International Conference on Composite Materials, Kyoto, Japan, 2007.
- [7] P. E. Rubbert, Cfd and the changing world of airplane design, in: 19th Congress of the International Council of the Aeronautical Sciences (ICAS), Anaheim, CA, 1994.
- [8] B. Grossman, Z. Gürdal, R. T. Haftka, G. J. Strauch, W. M. Eppard, Integrated aerodynamic/structural design of a sailplane wing, *Journal of Aircraft* 25 (9) (1988) 855–860. [doi:10.2514/3.45670](https://doi.org/10.2514/3.45670).
- [9] B. Grossman, R. T. Haftka, J. Sobieszczanski-Sobieski, P.-J. Kao, D. M. Polen, M. Rais-Rohani, Integrated aerodynamic-structural design of a transport wing, *Journal of Aircraft* 27 (12) (1990) 1050–1056. [doi:10.2514/3.45980](https://doi.org/10.2514/3.45980).
- [10] A. Chattopadhyay, N. Pagaldipti, A multidisciplinary optimization using semi-analytical sensitivity analysis procedure and multilevel decomposition, *Computers & Mathematics with Applications* 29 (7) (1995) 55–66. [doi:10.1016/0898-1221\(95\)00018-T](https://doi.org/10.1016/0898-1221(95)00018-T).

BIBLIOGRAPHY

- [11] M. Baker, J. Giesing, A practical approach to MDO and its application to an HSTC aircraft, in: 1st AIAA Aircraft Engineering, Technology, and Operations Congress, American Institute of Aeronautics and Astronautics, 1995, doi: 10.2514/6.1995-3885.
- [12] A. A. Giunta, V. Balabanov, D. Haim, B. Grossman, W. H. Mason, L. T. Watson, R. T. Haftka, Multidisciplinary optimization of a supersonic transport using design of experiments theory and response surface modeling, Tech. rep., Virginia Polytechnic Institute & State University, Blacksburg, VA, USA (1997).
- [13] K. Maute, M. Nikbay, C. Farhat, Coupled analytical sensitivity analysis and optimization of three-dimensional nonlinear aeroelastic systems, *AIAA Journal* 39 (11) (2001) 2051–2061. [doi:10.2514/2.1227](https://doi.org/10.2514/2.1227).
- [14] J. Reuther, J. Alonso, J. Martins, S. Smith, A coupled aero-structural optimization method for complete aircraft configurations, in: 37th Aerospace Sciences Meeting and Exhibit, American Institute of Aeronautics and Astronautics, 1999, doi: 10.2514/6.1999-187.
- [15] J. Martins, J. Alonso, J. Reuther, Aero-Structural Wing Design Optimization Using High-Fidelity Sensitivity Analysis, in: Proceedings of the CEAS Conference on Multidisciplinary Aircraft Design and Optimization, 2001.
- [16] J. Martins, J. Alonso, J. Reuther, High-fidelity aerostructural design optimization of a supersonic business jet, *Journal of Aircraft* 41 (3) (2004) 523–530. [doi:10.2514/1.11478](https://doi.org/10.2514/1.11478).
- [17] G. K. W. Kenway, G. J. Kennedy, J. R. R. A. Martins, Scalable parallel approach for high-fidelity steady-state aeroelastic analysis and adjoint derivative computations, *AIAA Journal* 52 (5) (2014) 935–951. [doi:10.2514/1.J052255](https://doi.org/10.2514/1.J052255).
- [18] G. K. W. Kenway, J. R. R. A. Martins, Multipoint high-fidelity aerostructural optimization of a transport aircraft configuration, *Journal of Aircraft* 51 (1) (2014) 144–160. [doi:10.2514/1.C032150](https://doi.org/10.2514/1.C032150).
- [19] G. Kennedy, J. Martins, A laminate parametrization technique for discrete ply-angle problems with manufacturing constraints, *Structural and Multidisciplinary Optimization* 48 (2) (2013) 379–393. [doi:10.1007/s00158-013-0906-9](https://doi.org/10.1007/s00158-013-0906-9).
- [20] T. Brooks, G. Kennedy, J. Martins, High-fidelity multipoint aerostructural optimization of a high aspect ratio tow-steered composite wing, in: 58th AIAA/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference, 2017.

-
- [21] E. Jonsson, G. Kenway, G. Kennedy, J. Martins, Development of flutter constraints for high-fidelity aerostructural optimization, in: 35th AIAA Applied Aerodynamics Conference, 2017. [doi:10.2514/6.2017-4455](https://doi.org/10.2514/6.2017-4455).
- [22] G. J. Kennedy, J. R. R. A. Martins, A parallel aerostructural optimization framework for aircraft design studies, *Structural and Multidisciplinary Optimization* 50 (6) (2014) 1079–1101. [doi:10.1007/s00158-014-1108-9](https://doi.org/10.1007/s00158-014-1108-9).
- [23] A. Elham, M. J. L. van Tooren, Coupled adjoint aerostructural wing optimization using quasi-three-dimensional aerodynamic analysis, *Structural and Multidisciplinary Optimization* 54 (4) (2016) 889–906. [doi:10.1007/s00158-016-1447-9](https://doi.org/10.1007/s00158-016-1447-9).
- [24] M. Barcelos, H. Bavestrello, K. Maute, A schur-newton-krylov solver for steady-state aeroelastic analysis and design sensitivity analysis, *Computer Methods in Applied Mechanics and Engineering* 195 (17-18) (2006) 2050–2069. [doi:10.1016/j.cma.2004.09.013](https://doi.org/10.1016/j.cma.2004.09.013).
- [25] M. Akgün, R. Haftka, C. Wu, J. Walsh, J. Garcelon, Efficient structural optimization for multiple load cases using adjoint sensitivities, *AIAA Journal* 39 (3) (2001) 511–516. [doi:10.2514/2.1336](https://doi.org/10.2514/2.1336).
- [26] N. Poon, J. Martins, An adaptive approach to constraint aggregation using adjoint sensitivity analysis, *Structural and Multidisciplinary Optimization* 34 (1) (2007) 61–73. [doi:10.1007/s00158-006-0061-7](https://doi.org/10.1007/s00158-006-0061-7).
- [27] A. Giunta, A novel sensitivity analysis method for high fidelity multidisciplinary optimization of aero-structural systems, in: 38th Aerospace Sciences Meeting and Exhibit, American Institute of Aeronautics and Astronautics, 2000, 10.2514/6.2000-683.
- [28] J. R. R. A., A coupled-adjoint method for high-fidelity aero-structural optimization, Ph.D. thesis, Stanford University (2002).
- [29] C. Blondeau, T. Achard, P. Girodroux-Lavigne, R. Ohayon, Recent achievements towards aero-structure gradient computation using high-fidelity cfd-csm in the onera elsa software, in: International Forum on Aeroelasticity and Structural Dynamics, 2015.
- [30] T. Achard, C. Blondeau, R. Ohayon, An uncoupled approach to compute aero-structure gradients using high-fidelity cfd-csm, in: 17th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference, 2016. [doi:10.2514/6.2016-4121](https://doi.org/10.2514/6.2016-4121).
- [31] G. Kreisselmeier, R. Steinhauser, Systematic control design by optimizing a vector performance index, *International Federation of Active Controls Symposium on Computer-Aided Design of Control Systems* 12 (7) (1979) 113–117.
-

BIBLIOGRAPHY

- [32] A. Lambe, G. Kennedy, J. Martins, An evaluation of constraint aggregation strategies for wing box mass minimization, *Structural and Multidisciplinary Optimization* 55 (1) (2017) 257–277. [doi:10.1007/s00158-016-1495-1](https://doi.org/10.1007/s00158-016-1495-1).
- [33] G. A. Wrenn, An indirect method for numerical optimization using the kreisselmeier-steinhauser function, Tech. rep., NASA Contractor Report 4220 (1989).
- [34] J. K. S. Dillinger, Static aeroelastic optimization of composite wings with variable stiffness laminates, Ph.D. thesis, Delft University of Technology, Delft (2014).
- [35] N. P. M. Werter, Aeroelastic modelling and design of variable stiffness composite and morphing wings, Ph.D. thesis, Delft University of Technology, Delft (2017).
- [36] N. P. M. Werter, R. De Breuker, A novel dynamic aeroelastic framework for aeroelastic tailoring and structural optimisation, *Composites Structures* 158 (2016) 369–386.
- [37] J. Katz, A. Plotkin, *Low-speed aerodynamics*, Cambridge University Press, 2001.
- [38] L. Cambier, M. Gazaix, S. Heib, S. Plot, M. Poinot, J. Veuillot, J. Boussuge, M. Montagnac, An overview of the multi-purpose elsA flow solver, *Aerospace Lab Journal* (2).
- [39] C. Hirsch, *Numerical Computation of Internal and External Flows, Vol. 2: Computational Methods for Inviscid and Viscous Flows*, Wiley, 1990.
- [40] MSC.Software Corporation, *MSC.Nastran 2012, Design Sensitivity and Optimization Users Guide* (2012).
- [41] MSC.Software Corporation, *MSC.Nastran 2013, DMAP Programmer's Guide* (2013).
- [42] A. Beckert, H. Wendland, Multivariate interpolation for fluid-structure-interaction problems using radial basis functions, *Aerospace Science and Technology* 5 (2) (2001) 125–134.
- [43] T. Rendall, C. Allen, Efficient mesh motion using radial basis functions with data reduction algorithms, *Journal of Computational Physics* 228 (17) (2009) 6231–6249.
- [44] T. Rendall, C. Allen, Reduced surface point selection options for efficient mesh deformation using radial basis functions, *Journal of Computational Physics* 229 (8) (2010) 2810–2820.

- [45] K. Jovanov, R. De Breuker, M. Abdalla, J. Dillinger, Accelerated convergence of static aeroelasticity using low-fidelity aerodynamics, in: 56th AIAA/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference, American Institute of Aeronautics and Astronautics, 2015.
- [46] P. Thévenaz, T. Blu, M. Unser, Interpolation revisited, *IEEE Transactions on Medical Imaging* 19 (7) (2000) 739–758.
- [47] R. Harder, R. Desmarais, Interpolation using surface splines, *Journal of Aircraft* 9 (2) (1972) 189–191.
- [48] J. Peter, R. Dwight, Numerical sensitivity analysis for aerodynamic optimization: A survey of approaches, *Computers & Fluids* 39 (3) (2010) 373–391. [doi:10.1016/j.compfluid.2009.09.013](https://doi.org/10.1016/j.compfluid.2009.09.013).
- [49] E. Nielsen, K. Anderson, Aerodynamic design optimization on unstructured meshes using the navier-stokes equations, *AIAA Journal* 37 (11) (1999) 1411–1419.
- [50] J. Peter, J. Mayeur, Improving accuracy and robustness of a discrete direct differentiation method and discrete adjoint method for aerodynamic shape optimization, in: *European Conference on Computational Fluid Dynamics, ECCOMAS CFD, Delft, Netherlands, 2006*.
- [51] Z. Luy, G. Kenway, C. Paige, J. Martins, Automatic differentiation adjoint of the reynolds-averaged navier-stokes equations with a turbulence model, in: *21st AIAA Computational Fluid Dynamics Conference, San Diego, CA, 2013*.
- [52] S. Yoon, A. Jameson, An lu-ssor scheme for the euler and navier-stokes equations, in: *AIAA 25th Aerospace Sciences Meeting, 1987*. [doi:10.2514/6.1987-600](https://doi.org/10.2514/6.1987-600).
- [53] Z. Zhang, Z. Zingg, Efficient monolithic solution algorithm for high-fidelity aerostructural analysis and optimization, *AIAA Journal* (2017) 1–15 [doi:10.2514/1.J056163](https://doi.org/10.2514/1.J056163).
- [54] A. Wathen, Preconditioning and convergence in the right norm, *International Journal of Computer Mathematics* 84 (8) (2007) 1199–1209.
- [55] A. Samareh, Application of quaternions for mesh deformation, *NASA TM-2002-211646*, 2002.
- [56] Y. Saad, *Iterative Methods for Sparse Linear Systems*, Society for Industrial and Applied Mathematics, 2003.
- [57] B. M. Irons, R. C. Tuck, A version of the aitken accelerator for computer iteration, *International Journal for Numerical Methods in Engineering* 1 (3) (1969) 275–277.

BIBLIOGRAPHY

- [58] X. Pinel, M. Montagnac, Block krylov methods to solve adjoint problems in aerodynamic design optimization, *AIAA Journal* 51 (9) (2013) 2183–2191. [doi:10.2514/1.J052113](https://doi.org/10.2514/1.J052113).
- [59] L. Giraud, S. Gratton, X. Pinel, X. Vasseur, Flexible gmres with deflated restarting, *SIAM Journal on Scientific Computing* 32 (4) (2010) 1858–1878. [doi:10.1137/080741847](https://doi.org/10.1137/080741847).
- [60] J. Walsh, Application of mathematical optimization procedures to a structural model of a large finite-element wing, Tech. rep., NASA TM-87597 (1986).
- [61] J. Sobieszczanski-Sobieski, Sensitivity of complex, internally coupled systems, *AIAA Journal* 28 (1) (1990) 153–160.
- [62] E. Anderson, Z. Bai, C. Bischof, S. Blackford, J. Demmel, J. Dongarra, J. Du Croz, A. Greenbaum, S. Hammarling, A. McKenney, D. Sorensen, *LAPACK Users Guide*, Society for Industrial and Applied Mathematics (1999).
- [63] G. Schrauf, Status and perspectives of laminar flow, *Aeronautical Journal* 109 (1102) (2005) 639–644.
- [64] M. Braun, K. Wicke, A. Koch, T. Wunderlich, Analysis of natural laminar flow aircraft based on airline network design and fleet assignment, in: 11th AIAA Aviation Technology, Integration, and Operations (ATIO) Conference, American Institute of Aeronautics and Astronautics, 2011. [doi:10.2514/6.2011-6807](https://doi.org/10.2514/6.2011-6807).
- [65] N. J. J. Krone, Divergence elimination with advanced composites, *AIAA Paper* (1975) 75–1009.
- [66] T. A. Weisshaar, Divergence of forward swept composite wings, *Journal of Aircraft* 17 (6) (1980) 442–448.
- [67] T. A. Weisshaar, Aeroelastic tailoring of forward swept composite wings, *Journal of Aircraft* 18 (8) (1981) 669–676.
- [68] L. Librescu, A. A. Khdeir, Aeroelastic divergence of swept-forward composite wings including warping restraint effect, *AIAA Journal* 26 (11) (1988) 1373–1377.
- [69] L. Librescu, J. Simovich, General formulation for the aeroelastic divergence of composite swept-forward wing structures, *Journal of Aircraft* 25 (4) (1988) 364–371.
- [70] L. Librescu, O. Song, On the static aeroelastic tailoring of composite aircraft swept wings modelled as thin-walled beam structures, *Composites Engineering* 2 (5–7) (1992) 497–512.

- [71] A. Seitz, M. Kruse, T. Wunderlich, J. Bold, L. Heinrich, The dlr project lamair: Design of a nlf forward swept wing for short and medium range transport aircraft, in: 29th AIAA Applied Aerodynamics Conference, Reston, VA, 2011.
- [72] K. Svanberg, A class of globally convergent optimization methods based on conservative convex separable approximations, *SIAM Journal on Optimization* 12 (2) (2002) 555–573. [doi:10.1137/S1052623499362822](https://doi.org/10.1137/S1052623499362822).
- [73] T. J. Hertz, M. H. Shirk, R. H. Ricketts, T. A. Weisshaar, Aeroelastic tailoring with composites applied to forward swept wings, Tech. rep., Flight Dynamics Laboratory (FIBR) AF Wright Aeronautical Laboratories, Wright-Patterson Air Force Base, Ohio 45433 (1981).
- [74] I. H. Abbott, A. E. Von Doenhoff, *Theory of Wing Sections: Including a Summary of Airfoil Data*, Dover Books on Aeronautical Engineering Series. Dover Publications, 1959.

BIBLIOGRAPHY

ACKNOWLEDGEMENTS

I believe that PhD students can be associated to one of two categories. The first category are those subject to an abundance of oversight and supervision from their supervisor, not affording the PhD students to think creatively in order to drive their research forward. These students, might have a good communication with their supervisor, but they are not exposed and encouraged to leave their comfort zone in order to develop their character and take responsibility for their research. The second category are the students on the other side of this spectrum. These students might not communicate on a regular basis with their supervisor, but are free to drive their research forward as they see fit. The responsibility lies on their shoulders and they are fully aware that the success (or failure) of their research depends on their ability to communicate with fellow researchers. Reflecting upon my journey as a PhD student, I definitely associate with the latter category. I am extremely grateful that my supervisor and promotor, dr. Roeland De Breuker, allowed me to take the reigns early on and shape my research into what it is today. He showed me great support from start to finish but never imposed his vision or ideas on my research. I always enjoyed his company and I consider him not only as a mentor, but as a very good friend.

I would also like to express my gratitude to Dr. Mostafa Abdalla who had a great impact on my research. I always enjoyed talking to one of the brightest minds, pitching ideas and getting feedback on the progress in this thesis. We could converse for hours about politics, culture and history. Concurrently, I would like to extend my deepest gratitude to monsieur Christophe Blondeau. To put it simply: without him, this PhD would not have been possible. He helped me modify the CFD solver elsA such that I would be able to use this powerful tool for my purposes in this thesis. I would like to thank Philippe Girodroux-Lavigne, the former head of the Aeroelasticity Department at ONERA, who hosted me at his department for several months and allowed me to conduct my research. I would also like to extend my gratitude to Andrea Viti and his supervisors at ONERA in Meudon for providing the consortium members of the Marie Curie project AMEDED with the CFD model of the forward-swept wing.

I would like to thank all my fellow PhD students, with whom I have so many fond memories. We traveled to conferences together, we had beers together, we played counter strike together and ordered Indian food... I will miss all of you. I would

BIBLIOGRAPHY

like to thank Laura Chant as well. She was the first one I met when I arrived at the office and I will never forget that day. She made me feel comfortable at my new setting and helped me a lot throughout my PhD. I would like to thank prof. Chiara Bisagni for being my promotor and giving me valuable feedback to improve my research. I would also like to thank the committee members, prof. Karpel, prof. Toropov, prof. Bijl, prof. van Bussel and prof. Benedictus for agreeing to be part my PhD defence committee.

Lastly, but most importantly, I would like to thank my family. My parents, Divna and Blaze, who never doubted my abilities to succeed, who always supported me and showed faith in me to become the man I am today. I am eternally grateful for your support and I could not have wished for more caring parents. I would also like to thank my big brother, Alexander. He brings out the best in me and has been an inspiration and role model of mine since we were kids. Finally I would like to dedicate this thesis to my grandmother, Trajana Dimova, who sadly passed away in March 2017. She will always have a special place in my heart.

BIOGRAPHICAL NOTE

Kristofer Jovanov was born on September 22, 1987 in Gothenburg, Sweden. He started studying Mechanical Engineering at Chalmers University of Technology in 2007. After obtaining his *Bachelor of Science* degree in 2010, he was awarded a scholarship and studied in Taiwan for one full year.

Upon his return from Taiwan he obtained the degree *Master of Science* by completing his thesis at a company, Creo Dynamics, in Lund, Sweden. The thesis, *Structural health monitoring for aerospace composite structures - an investigation of the potential using the finite element method*, was completed successfully and he was offered a job position in Lund as a FEM consultant, an offer he gladly accepted.



Kristofer continued working at the company in Lund, a city where his older brother resided in as well. His close bond with his brother helped him get accustomed to the new environment. Kristofer, however, is driven by knowledge and challenge. He desired more than a Master of Science degree. As soon as an opportunity presented itself in the form of a PhD position at TU Delft, he did not hesitate to send an application. One week after his first interview, he was already sitting on a plane heading towards the Amsterdam and ready to embark on a new adventure.

Upon finishing his PhD, he returned to Lund in Sweden. There he got hired by FS Dynamics, which is the leading consultant company in Sweden, specialized in performing FEM and CFD analyses for large industrial companies.

On his spare time Kristofer enjoys spending time with his family. He has an amazing brother, a lovely sister-in-law and an adorable niece and nephew. His caring parents live close by and have always supported him in good times and bad. He likes to work out in the gym and taking strolls in the park. He is not an active traveler, but enjoys visiting new places when an opportunity is presented.