Geometrically Nonlinear High Fidelity Aerostructural Optimisation for Highly Flexible Wings

MSc Thesis

Alasdair Christison Gray

TU Delft Delft University of Technology

MDO Lab

Geometrically Nonlinear High Fidelity Aerostructural Optimisation for Highly Flexible Wings

MSc Thesis

by

Alasdair Christison Gray

to obtain the degree of Master of Science

at the Delft University of Technology,

to be defended publicly on Tuesday August 11, 2020 at 2:00 PM.

Student number:	4737245	
Project duration:	September 10, 2019 – August 11, 2020	
Thesis committee:	Prof. Dr. ir. Leo Veldhuis,	TU Delft, Chair
	Dr. ir. Roeland De Breuker,	TU Delft, Supervisor
	Prof. Dr. Joaquim. R. R. A. Martins,	University of Michigan, Supervisor
	Dr. ing. Saullo. G. P. Castro,	TU Delft

An electronic version of this thesis is available at repository.tudelft.nl.

TU Delft Delft University of Technology

UNIVERSITY OF MICHIGAN

MDO Lab

Acknowledgements

4 years ago, I saw some videos on YouTube by a group called the 'MDO Lab', showing computer models of aircraft magically morphing themselves into better designs, and thought to myself that this would be a pretty cool thing to spend my time on. It is nice now to pause and take stock of the fact that, for the past year, I have been able to spend my time making my own versions. Despite all I have since learned about what goes on behind the scenes, the process still seems no less cool, and to some extent magical, than it did back then.

This thesis represents the culmination of 3 years of my life that have taken me from Edinburgh to Delft, several race tracks, Munich, Michigan and now, finally, in the midst of a global pandemic, my childhood bedroom, where I currently write these words. As with everything I do, it is the product not only of my individual efforts but also of the help given to me by many people over these past years. The very least I can do in return is thank them here:

My supervisors: To Quim Martins for welcoming me into your research group at Michigan, for treating me like any other member of the MDO Lab, for your support and guidance, and for letting me burn through a few thousand CPU hours of your cluster allowances. To Roeland De Breuker, firstly for inspiring my interest in aeroelasticity and secondly for agreeing to take on the remote supervision of this project. To Sandy Mader for sharing a fraction of his incredible technical intuition with me and asking just the right questions to point me in the right direction.

My MDO Lab colleagues: To Neil, John, Eirikur, Sham, Marco, Sabet, Shugo, Anil, Yinqian, Nick, Ben, Josh, Xiaoson, Ping, Anjali, Gustavo and Sicheng. For making me feel almost immediately at home in your group, for expanding my knowledge of computers 100-fold, for the lunches, for the dinners, and for the `hardcore_pasta`.

My housemates: To Feriel, Alana, Bre and Ian. For making home such a comfortable place to be, and for saying nice things about my cooking.

My DUT18 teammates: For a once in a lifetime experience on a rather time consuming side project. In particular to Till and Boris for making producing those 155 suspension design iterations a somewhat bearable process.

My friends in Delft: To Johannes, Tomas, Pablo, Zoe, Roberto, Filipo, Craig, Paula and all the others I shared my time in Delft with. For your friendship and support during a pretty crazy year, and for reminding me to look after myself when it seemed like I might be forgetting to.

Additional thanks go to Tim Brooks and Graeme Kennedy for their insight into the inner workings of TACS, to Paul Lancelot for his feedback during my mid-term review and to Leo Veldhuis and Saullo Castro for being on my defence committee, particularly to Saullo for stepping in at the 11^{th} hour.

Finally, to my family, and particularly my parents, Rona and Eric. Since leaving home 8 years ago, I have had the privilege of attending four different universities, working and studying in 4 different countries and building a couple of race cars along the way. Every one of these experiences has been enabled by the confidence I felt in pursuing them because of the support I had from you, both in emotional and more practical forms. Over the years I have become increasingly aware of how lucky I am to have this level of support and I cannot fully express how grateful I am for it. Saying all that, I'm sure you're both looking forward to me having my own income from now on.

Abstract

Over the past decade, advances in multidisciplinary design optimization (MDO) have enabled the optimization of aircraft wings using high-fidelity simulations of their coupled aerodynamic and structural behaviour. Using RANS CFD and detailed structural finite element wingbox models, the aerodynamic shape and internal structural sizing of a wing can be optimized concurrently, allowing the optimizer to correctly tailor the aeroelastic behaviour of the wing, resulting in the optimal trade-off between the competing goals of minimizing aerodynamic drag and structural mass. This capability makes MDO a key enabling technology for the next generation of efficient high-aspect-ratio transport aircraft. However, as their aspect-ratios increase, these wings increasingly exhibit geometrically nonlinear behaviour which cannot be correctly modelled by typical linear structural analysis methods. This thesis aims to bridge the gap between low-fidelity aeroelastic optimisation frameworks that include geometric nonlinearities and high-fidelity frameworks that cannot.

To do this, I extend the capabilities of the state of the art high-fidelity MDO framework, MACH, to include geometrically nonlinear analysis. I develop a Newton-Raphson based solver for the finite element library TACS which is integrated into MACH to facilitate geometrically nonlinear structural and coupled aerostructural analysis. The solver features an adaptive load incrementation strategy which includes a novel strain energy based method for restarting the load incrementation process from a previous solution. I also implement a load ramping strategy in MACH's partitioned aerostructural solver which I demonstrate can speedup a geometrically nonlinear aerostructural analysis by up to 40%.

I then perform comparisons of geometrically linear and nonlinear structural and aerostructural analysis on models of both a moderate and high aspect-ratio transport aircraft. The results show that, under the same loading, geometrically nonlinear analysis results in higher bending moments within the wingbox due to follower force effects, which in turn result in a 10% increase in the stresses in the wing skins in a 2.5 g pull-up manoeuvre. The nonlinear analysis also shows evidence of Brazier loading, a buckling critical compressive stress in the wingbox ribs that is not captured by linear analysis methods. In coupled aerostructural analysis, geometrically nonlinear span shortening effects are at the root of most of the differences between linear and nonlinear results. Most notably, a forward shift in the wing's centre of pressure changes the aircraft trim and the shortening of the effective wing span reduces the bending moment and stresses in the wingbox by around 5% when the aircraft is trimmed to the same 2.5 g load factor.

Finally I perform what are, as far as I am aware, the first published structural and aeroelastic optimisations of a geometrically nonlinear high-fidelity wingbox model. The structural optimisations to minimise the mass of the two wingboxes with respect to 6-700 structural sizing variables, subject to 15 failure constraints across 3 loadcases and 7-800 linear sizing constraints. In the structural optimisations the increase in bending due to geometric nonlinearity lead to a 6% increase in the mass of the optimised high aspect-ratio wingbox, and a smaller but still significant 4% increase in mass in the moderate aspect-ratio model. Brazier loads also lead to significant changes in the sizing of the wing ribs but a negligible increase in their mass.

The results of the aeroelastic optimisations appear to contradict the observed decrease in stress seen in the initial aerostructural analysis comparisons with the geometrically nonlinear optimised design being 3% heavier than the same wing optimised with geometrically linear analysis. It appears that this discrepancy may be caused by erroneous stress values in distorted structural mesh elements and is a topic that merits further investigation. Encouragingly, the increase in the computational cost of these high-fidelity aeroelastic optimisations due to the addition of geometrically nonlinear structural analysis appears to be less than 20%.

Nomenclature

General Mathematical Notation

$\mathcal{F}$	Function of Interest
$\mathcal{J}$	Jacobian
$\mathcal{O}$	Order of Magnitude
$\mathcal{R}$	Residual
$\left[\mathcal{I}\right]$	Identity Matrix
$\left[A\right]$	Matrix A
ψ	Adjoint Variable
$\left\{A\right\}$	Column Vector A
u	State Variable
x	Design Variable

Finite Element and FSI Notation

α	Line SearchStep Length
δ_{abs}	Absolute Convergence Tolerance
δ_{rel}	Relative Convergence Tolerance
ϵ	Strain
λ	Load Factor
λ_{stiff}	Stiffener Pitch
$\left[K_T\right]$	Tangent Stiffness Matrix
$\left[K\right]$	Stiffness Matrix
$\left[T\right]$	Load-Displacement Transfer Matrix
θ	Under-relaxation Factor
$\left\{u_A\right\}$	Aerodynamic Surface Displacements
$\left\{w\right\}, \left\{v\right\}$	BFGS Update Vectors
$\left\{X_J\right\}$	Jig Aerodynamic Surface Coordinates
$\left\{X_S\right\}$	Displaced Aerodynamic Surface Coordinates
E	Strain Energy
F_s	Fluid Forces on Structural Mesh
F_{ex}	External Force
F_{in}	Internal Force
$F_{x,cr}$	Critical Compressive Load
$F_{xy,cr}$	Critical Shear Load
h_{stiff}	Stiffener Height
T_F	Force Transfer Process
T_X	Displacement Transfer Process
t_{skin}	Skin Thickness
t_{stiff}	Stiffener Thickness
X_f	Structural Displacements of Fluid Mesh

Other Symbols

α	Angle of Attack
θ_{tail}	Tail Rotation Angle
C_D	Drag Coefficient
C_L	Lift Coefficient
C_M	Moment Coefficient
COP	Centre of Pressure
$MTOW$	Maximum Take-off Weight

Acronyms

AD Algorithmic Differentiation. , 16, 20, 102

CFD Computational Fluid Dynamics. , 2, 4–6, 10, 11, 13, 16, 20, 37, 38, 42, 45, 46, 52, 53, 55, 56, 62, 78, 83, 85, 88, 90, 99, 101, 103

COG Centre of Gravity. , 53

COP Centre of Pressure. , 53, 55, 56, 59, 60, 86, 88, *Glossary:* centre of pressure

CRM Common Research Model. , 42

DLM Doublet Lattice Method. , 3, 9, 13

FE Finite Element. , 3–6, 9, 10, 13, 15, 16, 18, 22, 24, 28, 30, 37, 38, 45, 46, 52, 53, 55, 56, 78, 79, 85, 101, 104

FFD Free-form Deformation. , 21, 53

FSI Fluid-Structure Interaction. , 15, 36

HALE High Altitude, Long Endurance. , 89, 90

HARW High-aspect-ratio wing. , 1, 2, 4, 6, 13–15, 23, 62, 83, 85, 88, 89

KS Kreiselmeier-Steinhauser. , 4, 18, 63, *Glossary:* Kreiselmeier-Steinhauser function

MAC Mean Aerodynamic Chord. , 42

MACH MDO of Aircraft Configurations at High-fidelity. , 4, 5, 13, 15, 16, 18–20, 22, 37, 41, 78, 85, 86, 89, 90

MDO Multidisciplinary Design Optimisation. , 1, 2, 4, 5, 85, 99, 100, 102

MTOW Maximum Take-off Weight. , 44, 78

OML Outer Mould Line. , 20

RANS Reynolds-averaged Navier-Stokes. , 2, 4, 5, 11, 13, 16, 20, 78, 85

RLT Rigid Link Transfer. , 19, 20, 90

ROM Reduced-order Model. , 2, *Glossary:* reduced-order model

SQP Sequential Quadratic Programming. , 63

SVD Singular Value Decomposition. , 6

TACS Toolkit for Analysis of Composite Structures. , 4, 14, 15, 18–23, 29–31, 41, 42, 45, 78, 85, 87–90, 104

uCRM Undeflected Common Research Model. , xiii, 14, 15, 33, 36, 42–44, 46–48, 51–57, 59, 60, 62, 64, 68–71, 76–79, 82–84, 86, 87, 89, 90

VLM Vortex Lattice Method. , 5, 6, 9, 10, 12, 13

XDSM eXtended Design Structure Matrix. , 16

Glossary

aeroelastic optimisation Optimisation of a wing's structural sizing only, albeit using coupled aeroelastic analysis.. , 5

aerostructural optimisation Simultaneous optimisation of the aerodynamic shape and structural sizing of a wing using coupled aeroelastic analysis..

centre of pressure A point in space at which the total aerodynamic forces acting on a surface can be represented as a single resultant force vector..

Kreiselmeier-Steinhauser function A function which generates a smooth approximation of the maximum value of a set of functions f_i:

$$KS\left(f_i\left(x\right)\right) = \frac{1}{\rho} \ln\left(\sum_{i}^{N_f} e^{\rho f_i(x)}\right)$$

.

reduced-order model A simplification of a high-fidelity model that preserves essential behaviour and dominant effects, for the purpose of reducing the solution time or storage capacity required for the more complex model [1].

[1] Source:ANSYS, 2017, Reduced-Order Modelling How and Why An Overview

List of Figures

List of Tables

Contents

1

Introduction

Whether for environmental or economic reasons, aircraft manufacturers are constantly striving to design and build more efficient aircraft. One of the conceptually simplest ways to increase the aerodynamic efficiency of aircraft is to increase the aspect-ratio of their wings. Consequently, the past five decades have seen a steady increase in the aspect-ratio of commercial aircraft wings.

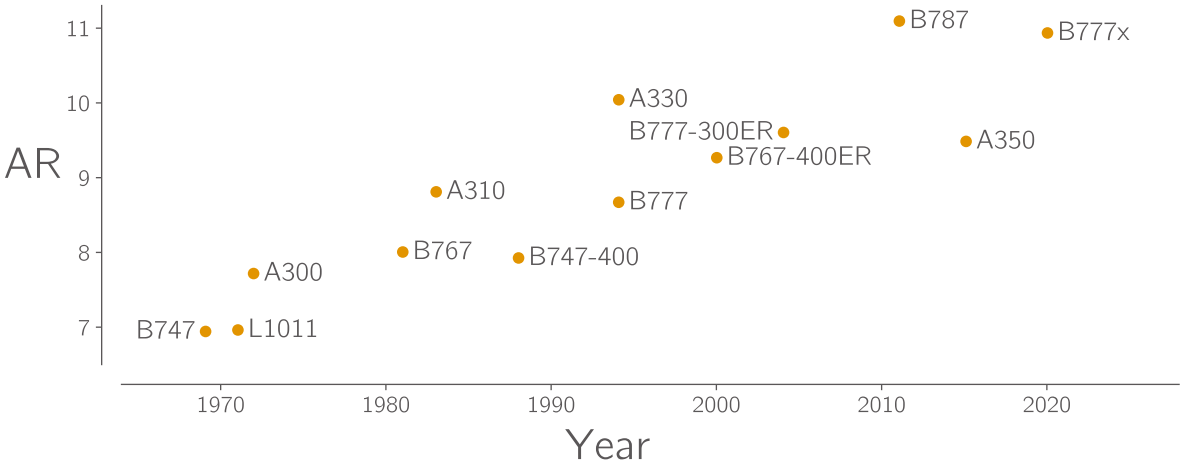

Figure 1.1: Commercial transport aircraft wings are increasing in aspect-ratio [1].

These developments are enabled by two distinct areas of innovation. First, the development of novel airframe technologies such as advanced composite materials, active load alleviation and aeroelastic tailoring of wing structures, and second, the development of new design techniques to extract the full potential of these new airframe technologies.

The need for new design methodologies for high-aspect-ratio wings (HARW) stems from two separate problems. The first is that as the aspect-ratio of wings grow, they naturally become more flexible, leading to a stronger coupling between their structural and aerodynamic behaviour and thus requiring coupled analysis techniques. Secondly, many of the novel airframe technologies mentioned above, which are key enablers of HARW, bring with them an increased number of design parameters, such as ply stacking sequences for composites or control surface deflections for active load alleviation, further complicating the already complex task of wing design.

In response to this need, there has been increased interest in the application of multidisciplinary design optimisation (MDO) methods to the design of aircraft, and in particular to aerostructural wing design. MDO tackles the two design challenges just mentioned by utilizing multidisciplinary analysis to fully capture coupled aerostructural behaviour and by using gradient based optimization techniques to efficiently search

complex, high dimensional design spaces. In the past decade, with increases in computational power it has become possible to perform the concurrent optimization of the aerodynamic shape and structural sizing of wings with high fidelity, using Reynolds-averaged Navier-Stokes (RANS) computational fluid dynamics (CFD) and detailed wingbox structural finite element models. Such methods have been used to perform aerostructural optimization of transonic commercial aircraft wings of both traditional design and utilizing novel technologies such as tow steered composites and wing morphing [2–6]. The end goal of these MDO techniques is to allow for as large a portion of the aircraft design process as possible to be based on accurate physics based simulations rather than on empirical data, the lack of which is currently the largest barrier to the development of novel aircraft.

Due to this replacement of engineering intuition with simulation results, the computational methods used in MDO frameworks must be capable of modelling the physical phenomena that limit the design space in reality in order to avoid unrealistic optimal designs. Moreover, the strong interdisciplinary couplings inherent to MDO problems means that modelling inaccuracies in one discipline lead to design inaccuracies in many others. This presents a problem for many current aerostructural MDO frameworks, which use linear structural analysis, as HARW exhibit geometrically nonlinear behaviour due to their large displacements.

Geometrically nonlinear finite element methods for structural analysis are a mature technology available in most commercial structural analysis codes and have been incorporated into a number of aeroelastic analysis frameworks in the past two decades [7–20]. More recently, these frameworks have been extend to perform geometrically nonlinear aeroelastic optimisation[1] of HARW [11, 21–23]. The scope of these optimisation problems is however limited to the sizing of the wing structure to minimise mass, excluding any geometric parameterisation of the wing. This is because the panel based aerodynamic models used in these frameworks do not allow for the accurate prediction of drag, especially in the transonic cruise conditions of most commercial aircraft.

The commonality between all of these geometrically nonlinear approaches is the use of lower fidelity structural models, either beam finite element formulations or reduced-order models (ROM) of full wingbox[2] finite element models. Although these simplified models have been proven capable of accurately predicting the global deflected shape of a wing, they are not able to predict detailed stress distributions within the wingbox components due the condensation inherent in their formulation. There is therefore a knowledge gap in the detailed aerostructural design of HARW resulting from a capability gap between low-fidelity aeroelastic optimisation frameworks that include geometric nonlinearities and high-fidelity aeroelastic optimisation frameworks which cannot.

In this chapter I first review the history and current state of the art in aerostructural design optimisation to further define the knowledge gap this thesis aims to address. I then review the basics of geometrically nonlinear structural phenomena before summarising the literature on how these geometrically nonlinear effects influence the aeroelastic performance and design of highly flexible wings. Finally, I restate the research problem based on the reviewed literature before forming the aims and research questions of this thesis.

1.1. Background

1.1.1. Aerostructural Wing Optimisation

In the 1960's and 70's interest in design optimisation grew rapidly in the academic aerospace community although, as described by Ashley [24], most applications were restricted to a single discipline and few, if any of these techniques were used in industry.

[1] I reserve the term 'aerostructural' for optimisation problems involving both the structural and aerodynamic design of a wing, whilst aeroelastic optimisation considers only structural design variables.

[2] Throughout this work, I use the terms 'full', 'detailed' 'shell' and 'high-fidelity' to refer to structural models that explicitly model the 3D geometry (e.g skins, spars, ribs) of a wingbox using shell elements, as opposed to models which condense wingbox sectional properties, as is done with beam finite elements.

One of the first applications of optimisation or 'nonlinear programming' to multidisciplinary design problems was in a tool developed by General Dynamics for the US Airforce in the early 1970's known as TSO (aeroelastic Tailoring and Structural Optimization) [25, 26]. The code used a Rayleigh-Ritz equivalent plate structural model and a doublet lattice method (DLM) aerodynamic code to optimise laminate orientations, thickness distributions and ballast masses subject to constraints on strength, and flutter velocity.

In 1977 Haftka built on his background in structural optimisation by coupling a panel based aerodynamics model with a basic wingbox finite element (FE) model to perform aeroelastic analysis and optimisation of a wing structure subject to stress and drag constraints [27]. The code was used to perform wing weight vs drag tradeoffs for aluminium and composite wings. Unsurprisingly, Haftka found that the composite wings were lighter and had lower drag (due to the decreased lift required) but that their increased flexibility resulted in a larger variation in drag along the pareto front.

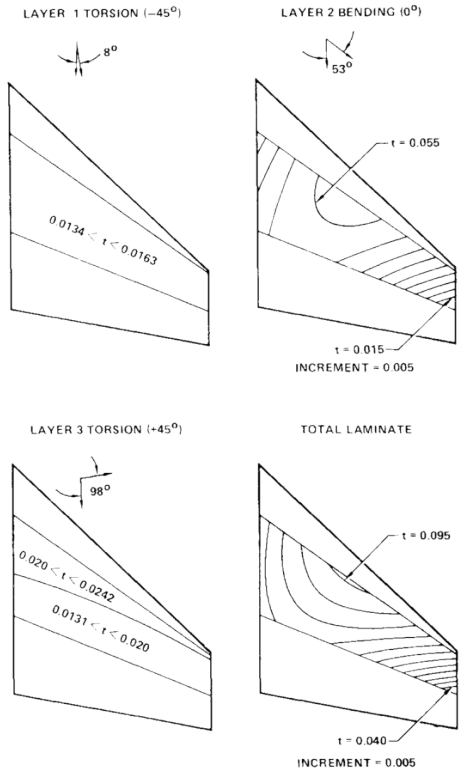

(a) Optimised laminate orientations and thicknesses distributions generated by TSO [25]

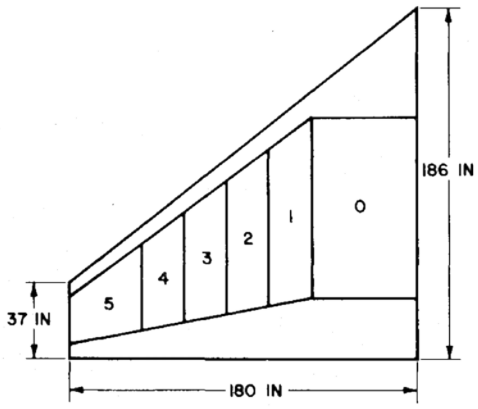

(b) Haftka's wing model [27]

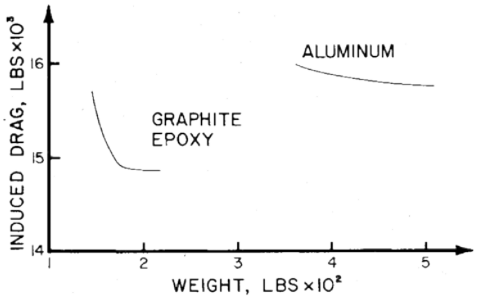

(c) Haftka's aerostructural pareto fronts [27]

Figure 1.2: The early days of aerostructural design optimisation

In some sense, neither of these works were truly aerostructural as only they concerned only the structural design of the wing, albeit incorporating aeroelastic interactions. Grossman et al. [28] performed one of the first examples of truly aerostructural optimisation. Using a lifting line aerodynamics model and analytical beam equations, they optimised the shape and structural sizing of a sailplane wing. Importantly, they demonstrated that a sequential aerodynamic and structural optimisation approach yielded worse performing wing designs than the 'integrated' approach. The integrated design approach was also better able to take advantage of beneficial passive aeroelastic effects, with the authors noting that: 'The integrated designs were characterized by less rigid, higher-aspect ratio wings that utilized favourable aerodynamic/structural interactions.'.

Later, Grossman et al. [29] performed a similar optimisation of a transport aircraft wing, this time using more complex panel aerodynamics and FE wingbox models. The main focus of that work was on an efficient method for calculating the derivatives of objective and constraint functions which depended on these complex multidisciplinary analyses. Indeed, around this time, a number of authors began focusing

their work on the efficient and accurate calculation of these derivatives [30] as they were increasingly becoming the bottleneck in both structural and aerostructural design optimisation processes, as Adelman and Haftka [31] explained:

> Early attempts to use formal optimization for large structural systems resulted in excessively long and expensive computer runs. Examination of the optimization procedures indicated that the predominant contributor to the cost and time was the calculation of derivatives.

The subject of gradient computation is of great importance in MDO, particularly when using high-fidelity models. I do not cover the ins and outs of these techniques in this chapter as it not of relevance to this thesis's key contributions but a comprehensive overview is given in appendix A. Save to say that when computing derivatives of functions which depend on the solution of large systems of equations, there are two primary methods to efficiently obtain accurate results. The direct method, which can efficiently compute the derivative of many outputs with respect to a few inputs, and the adjoint method, which can efficiently compute the derivatives of a few outputs with respect to many inputs.

With growing computational power, came a corresponding increase in the fidelity of aerodynamic and structural models used in aerostructural optimisation. Maute presented some of the first examples of aerostructural optimisations using 3D Euler computational fluid dynamics (CFD). Maute et al. [32] presented a 'three-field formulation' that coupled the solutions of the aerodynamic and structural models to another system of equations describing the deformation of the CFD mesh. They presented both direct and adjoint methods for calculating derivatives of the coupled system and demonstrated the power of the adjoint method by maximising the lift to drag ratio of a wing using 432 design variables controlling composite laminate orientations in a plate representing the wing structure. Later the same authors demonstrated an optimisation of both wing shape and structure using a realistic wingbox structure [33]. However, due to a large number of stress constraints (one at each node of the FE mesh), the authors used the direct method for calculating derivatives and were thus limited to using 5 design variables. By aggregating the stress values into a single constraint using a Kreiselmeier-Steinhauser (KS) function Maute et al. [34] were able to perform a similar optimisation with 56 design variables. Barcelos built on Maute's framework, developing a more robust Newton-Krylov-Schur approach for solving the coupled system and extending the framework to work with Reynolds-averaged Navier-Stokes (RANS) CFD [35, 36]. However, this work was limited to direct derivative calculation methods and thus few design variables. Barcelos cites issues developing an efficient parallel implementation for some aspects of the adjoint derivative calculation as the reason for this [36].

Around the same time, Martins et al. [37, 38] also developed a coupled adjoint for a coupled Euler CFD and linear FE model and demonstrated its ability to calculate derivatives with respect to $\mathcal{O}\left(10^3\right)$ variables. They used their framework to perform the aerostructural optimisation of a supersonic business jet with respect to 97 wing shape and sizing design variables [39].

In the past decade, the MDO Lab, headed by Martins, has developed the MACH (MDO of Aircraft Configurations at High-fidelity) framework which represents the current state of the art in high-fidelity aerostructural optimisation. The framework includes CFD (ADflow) and FE (TACS (toolkit for analysis of composite structures)) solvers with efficient adjoint derivative implementations along with fully differentiated modules required for coupled analysis and optimisation (e.g. geometry parameterisation, load and displacement transfer, and mesh warping) [6, 40–43]. Using MACH, Kenway, Kennedy and Martins were able to perform aerostructural optimisation at a new scale, with $\mathcal{O}\left(10^3\right)$ design variables and $\mathcal{O}\left(10^6 - 10^7\right)$ state variables, first with Euler and then RANS CFD [5, 44]. Kennedy and Martins [44] showed that allowing the optimiser to control the wing's airfoil profiles whilst using Euler CFD lead to unrealistically thick wings designs due to the Euler model's inability to predict shock induced boundary layer separation, thus demonstrating the importance of using a RANS aerodynamic model in enabling an accurate aerostructural trade-off. glsmach has since been used to investigate the potential benefits of some of the new airframe technologies enabling HARW. Burdette studied the effect of wing morphing in various works [3, 45–49] and Brooks extended the capabilities of TACS in order to study the performance benefits of tow-steered composite wings [2, 50–54]

Some alternative lower fidelity aerostructural optimisation frameworks include OpenAeroStruct developed

by Jasa et al [55, 56] and FEMWET, developed by Elham and van Tooren [57]. OpenAeroStruct uses the common combination of vortex lattice method (VLM) and beam FE models along with empirical corrections for viscous and wave drag. Although primarily used for educational purposes, it has been shown to produce similar results to higher fidelity frameworks in some cases [58]. FEMWET couples a 'quasi 3D' aerodynamic model to a linear beam model. The 'quasi 3D' aerodynamic model uses the VLM code AVL to compute a lift distribution which provides the lift coefficients for a series of 2D sectional analyses in a higher fidelity aerodynamic tool such as MSES or XFoil is used to compute parasitic drag. The inclusion of viscous and compressible flow effects in these sectional analyses allows FEMWET to accurately model the trade-off between drag and structural mass as is necessary for accurate aerostructural optimisation.

Aeroelastic Optimisation, an Alternate Path

Venkatamaran and Haftka [59] described the complexity of structural design optimisation as having 3 axes:

Model complexity: How detailed is the model used? (e.g what kind of elements and how many of them?)

Analysis complexity: How complex is the analysis performed? (e.g linear vs nonlinear, static vs dynamic vs spectral)

Optimisation complexity: How complex is the optimisation problem? (How many loadcases, how many variables etc)

They concluded that any given design optimisation problem can achieve maximum complexity in at most two of these three measures.

The MACH framework contains elements of all these types of complexity. With model complexity coming from the use of 3D CFD and full wingbox models, analysis complexity from the use of the RANS equations for aerodynamic analysis, and optimisation complexity due the large number of design variables used. However, in achieving this it also sacrifices other types of analysis and optimisation complexity. Optimisation complexity suffers as only a small ($\mathcal{O}(1-10)$) number of loadcases can be considered and analysis complexity is limited by the use of only steady state analysis and linear structural analysis.

A parallel trend has been in the development of aeroelastic optimisation frameworks which use lower fidelity panel aerodynamic models and often, though not always, beam FE models. These frameworks cannot be used for true aerostructural optimisation as they cannot accurately predict the transonic and viscous effects that are critical for drag prediction and thus cannot accurately trade-off drag and structural mass. They are therefore typically limited to minimising structural mass for a fixed wing shape. By reducing model complexity however, these frameworks allow for gains in analysis and optimisation complexity, some examples of which include:

Transient loadcases: Due to the lower cost of these aeroelastic analyses, these frameworks can perform aeroelastic analysis considering transient loadcases such as gust encounters. For example, Rajpal et al. [60] used the PROTEUS framework, developed by Werter and De Breuker [21], to perform aeroelastic optimisation considering around 20,000 different gust encounter simulations. Stanford [61] optimised a full wingbox model using a stochastic gust analysis technique proposed by Engelsen and Livne [62].

Fatigue life constraints: The ability to consider such large numbers of loadcases also allows for a more realistic consideration of fatigue life as a structural constraint. Rajpal et al. [63] developed a fatigue life failure analysis method for PROTEUS based on the linear wear-out law and the Tsai-Wu failure criterion. They found that using this fatigue prediction resulted in lighter wings than those where fatigue was accounted for with a simple knock down factor on strength.

Inclusion of control discipline: Panel based aerodynamic models make the modelling of control surfaces significantly easier than with CFD models, allowing these lower fidelity aeroelastic frameworks to include the design of control laws in the MDO problem. For example, as well as optimising the wingbox structure, Stanford [61] simultaneously optimised control surface laws for manoeuvre and gust load alleviation.

Complex Structural Parameterisations: Using high fidelity structural models coupled to low fidelity aerodynamics models, has allowed Stanford and collaborators to use more complex structural parameterisations which would otherwise be impractical for higher fidelity studies. These include topology optimisation of wingbox panels [64, 65], optimisation of the number and location of ribs and stiffeners [66] and curvilinear ribs and stiffeners [67].

Geometrically nonlinear structural models: The UM/NAST framework developed by Cesnik et al and PROTEUS both use geometrically nonlinear beam models. Lupp and Cesnik [19] used UM/NAST to perform basic aerostructural optimisation with a geometrically nonlinear flutter constraint whilst PROTEUS has been used for many aeroelastic optimisation studies [21, 23, 60, 63, 68–70].

Many works in recent years have focused on developing multifidelity methods, which use lower fidelity models to speed up the convergence of high fidelity models. Dillinger et al. [71] created an aeroelastic optimisation framework which combined low fidelity linear aeroelastic analysis in NASTRAN with intermittent load corrections from Euler CFD analysis in order to converge to the same solution as a standard high fidelity coupled analysis with lower cost. Jovanov [72] extended this method of high fidelity corrections to the computation of coupled derivatives in PROTEUS. Their method computes a high fidelity correction to the VLM aerodynamic stiffness matrix using an energy based singular value decomposition (SVD) of the Euler model's aerodynamic stiffness matrix. Aeroelastic optimisations performed with these high fidelity corrections applied to both analysis and gradients converged to a lower mass than cases using only a VLM model or only applying corrections to the analyses.

1.1.2. Geometrically Nonlinear Structural Phenomena

Linear structural analysis methods rely on the assumption that the displacements, strains, stresses and forces in a structure are all linearly related to one another. For many engineering applications, where structures should remain in the linear elastic region and undergo small displacements, this is a good assumption and simplifies any analysis to the solution of a single system of linear equations. There are however multiple scenarios in which these assumptions do not hold:

Material nonlinearity: Structures which see large stresses and strains may exhibit nonlinear material behaviour which can be elastic (e.g. hyperelasticity) or inelastic (e.g. yielding and plastic deformation). Modelling this behaviour requires a nonlinear stress-strain relationship.

Contact nonlinearity: The initiation of contact between two surfaces introduces strong discontinuous nonlinearities. Typically in FE problems, contact is modelled as an external force which has a nonlinear and discontinuous relationship with nodal displacement.

Geometric nonlinearity: When a structure undergoes large displacements, especially large rotations, the change in its shape can have a non-negligible effect on both the magnitude and direction of it's stiffness and the applied loads. Modelling these effects requires nonlinear displacement-strain and external force-displacement relationships respectively.

Since a typical wing structure does not see any contact phenomena and we are not concerned with modelling any post-failure structural behaviour, both contact and material nonlinearity are not considered further in this work. Geometric nonlinearity, however, is important for the structural analysis of high-aspect-ratio wings (HARW).

Figure 1.3, shows qualitative examples of the three dimensional stiffness distributions of a point on a cube and a shell visualised as ellipsoids. The stiffness of the point on a cube is roughly equal in all directions, resulting in a roughly spherical stiffness ellipsoid, in contrast the shell's significantly lower out-of-plane stiffness results in a flattened ellipsoid. Due to the rotational symmetry of a sphere, rotating the cube has little effect on the stiffness distribution in the global coordinate system. For the shell however, out-of-plane rotations cause a considerable change in the orientation of the principal stiffness directions and, consequently, the stiffness distribution as viewed in the global coordinate system. This behaviour is one reason that shell structures, such as wingboxes, exhibit strong geometric nonlinearity.

Additional large rotation nonlinearities appear when using follower forces, external loads which follow the

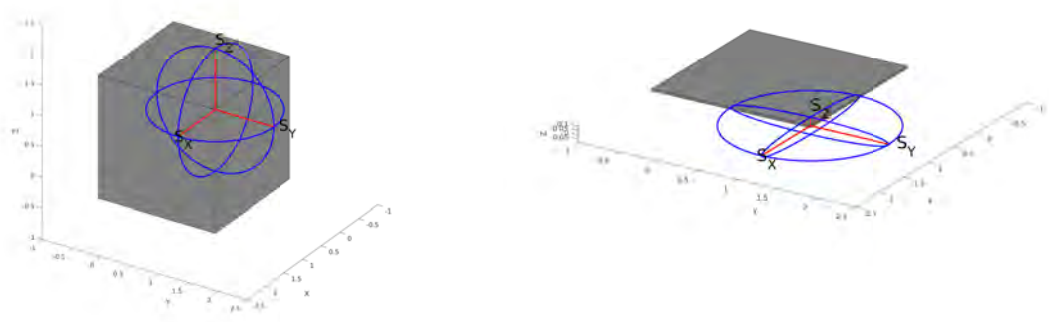

(a) Cube (b) Shell

Figure 1.3: Qualitative examples of the directional stiffness distributions for a cube and a shell element. The large difference between the in and out-of-plane stiffness of the shell element make its behaviour more sensitive to rotations than the cube.

rotation of the structure they are applied to. Aerodynamic loads can be considered follower forces as they consist of pressure and shear forces which remain normal and parallel to the wing surface respectively.

Figure 1.4 shows an example of these large rotation nonlinearities in the analysis of a cantilever beam subject to a uniform pressure load which exhibits follower force effects. In the linear model, the orientation of both the forces and of the structure's stiffness are only considered in its undeformed state. As a result, the model shows purely vertical deformation, resulting in an unphysical stretching of the beam. The geometrically nonlinear model correctly retains its original length and bends back on itself due to the rotation of pressure load.

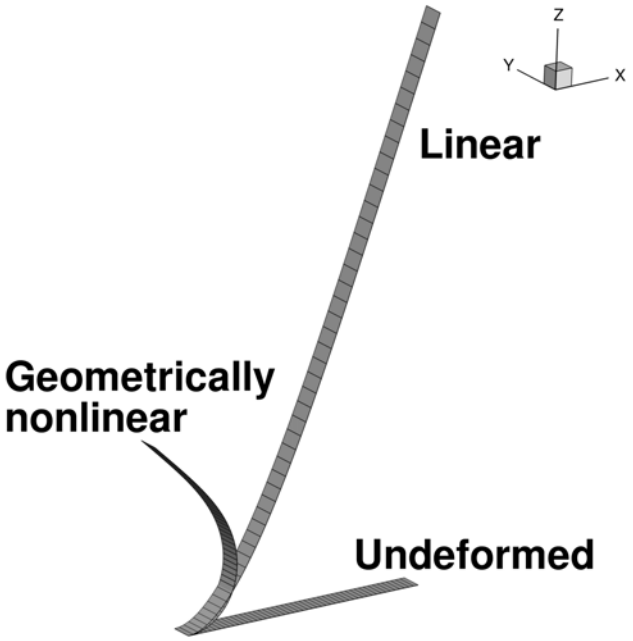

Figure 1.4: Cantilever with uniformly distributed pressure load which follows rotation. The linear model produces highly unrealistic results due to the large rotations present.

Another geometrically nonlinear effect that stems from the large discrepancy between in and out-of-plane stiffness of shells is stress stiffening. When highly loaded, the forces resulting from in-plane, or membrane,

stresses are great enough that they contribute to out-of-plane stiffness when out-of-plane deformations causes them to become oriented with the out-of-plane external forces. Tensile membrane stresses lead to increase in out-of-plane stiffness whilst compressive membrane stress causes decrease in out-of-plane stiffness which can eventually cause buckling.

Figures 1.5 and 1.6 show examples of these two forms of stress stiffening. Figure 1.5 shows a simply supported plate, again subject to uniform pressure load. In the nonlinear case, as it deflects downwards, significant tension is induced in the plate which, due to the deformation of the plate, resists the vertical loads. The linear case does not capture this stress stiffening effect as it does not consider the plate's deformed shape. Only the much lower bending stiffness of the plate is considered, resulting in unrealistically large deformations.

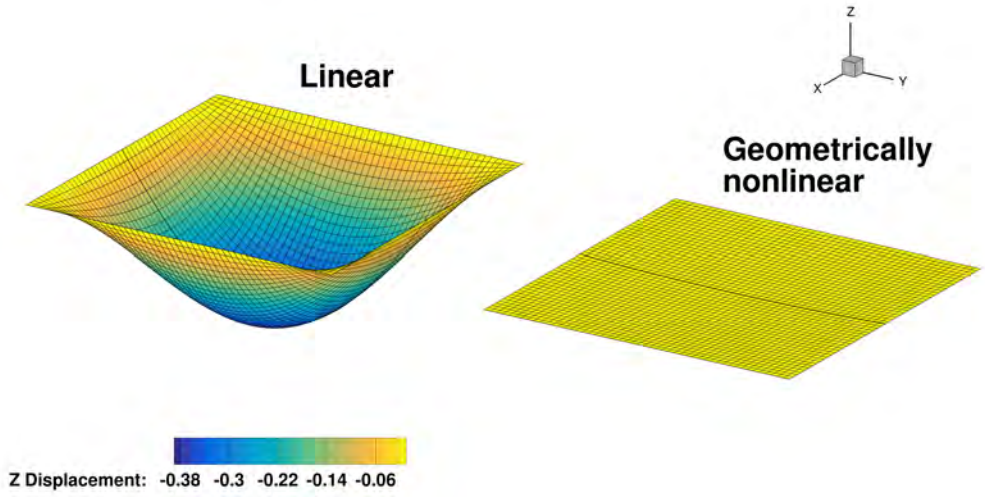

Figure 1.5: Plate subject to a uniform pressure load, linear model shows significantly larger deformation due to absence of stress stiffening effects.

The final example in figure 1.6 shows a thin walled box section beam subject to a vertical tip load. This case demonstrates the opposite kind of stress stiffening. As the upper skin is placed in compression, stress stiffening causes its out-of-plane stiffness to reduce to zero, resulting in buckling.

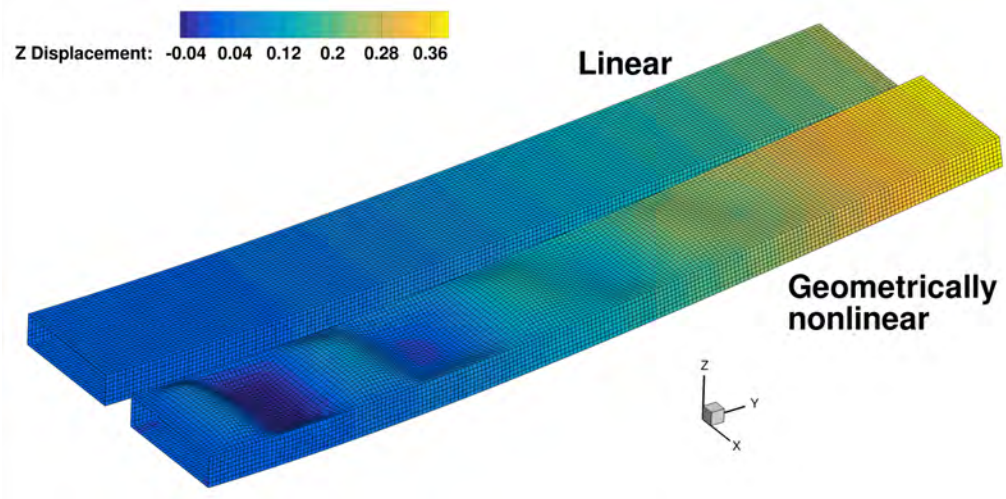

Figure 1.6: A thin walled box section cantilever beam subject to a tip load (deformations exaggerated 10×). The geometrically nonlinear model is able to capture buckling of the upper skin.

In thin walled beams subject to bending, stress stiffening due to the compressive and tensile stresses in the upper and lower surfaces of the beam cause a flattening of its cross section. This effect was discovered

by engineers building early aircraft wings from thin steel shells and is named after Brazier [73] who first published work on phenomena and showed how it can lead to the buckling of a beam's cross section. In a wingbox, Brazier effects lead to large compressive loads in the wing's ribs. These Brazier loads are typically the critical loads in the structural sizing of ribs, particularly for buckling and consequently, structural optimisation of wingboxes using linear FE models results in unrealistically lightweight rib designs [74]

1.1.3. Geometrically Nonlinear Effects in Aerostructural Analysis and Design

To correctly model aeroelastic phenomena under large deformation, both the structural and aerodynamic models used must contain geometric nonlinearities. The motivation for geometrically nonlinear structural models has just been given, but it is equally important that the aerodynamic model used correctly represents the geometry of the deformed wing in order to correctly model the changes in the magnitude and orientation of aerodynamic forces that occur under large displacements and rotations.

Linear VLM and DLM aerodynamic models do not explicitly model the twisting of wing sections, instead adding linearised terms to the non-penetration condition to account for small angles of attack of each panel. As a result, these models not only incorrectly predict the magnitude of aerodynamic loads under large rotations but also incorrectly model the aerodynamic forces as acting perpendicular to the undeformed aerodynamic surface.

Howcroft et al. [15] demonstrated the inaccuracy of this approach in a paper comparing a variety of different linear and nonlinear aeroelastic models including the academic nonlinear code NeoCASS [11] and both linear and nonlinear modes from the widely used commercial code Nastran. Figure 1.7a shows the lift (in the global vertical direction) distributions of linear and nonlinear aeroelastic models of a highly-flexible wing, trimmed to the same total lift level. The substantial difference in the distributions is due to the fact that the nonlinear model correctly captures the reduction in useful lift produced by the outboard wing sections due to the inboard rotation of the lifting surface. As a result, the nonlinear model requires a higher angle of attack to achieve the same lift and more of the useful lift is produced by the inboard wing sections. Figure 1.7b demonstrates the same effect from another perspective, showing the total lift produced by a linear and nonlinear aeroelastic model relative to that produced by a rigid version of the same wing over a range of flight speeds at a constant angle of attack [75]. At low speeds, both flexible models generate more lift than the rigid model due to some wash in at the wingtip. However, with increasing flight speed, the inboard rotation of the lift in the nonlinear model results in a loss in total lift relative to the rigid wing.

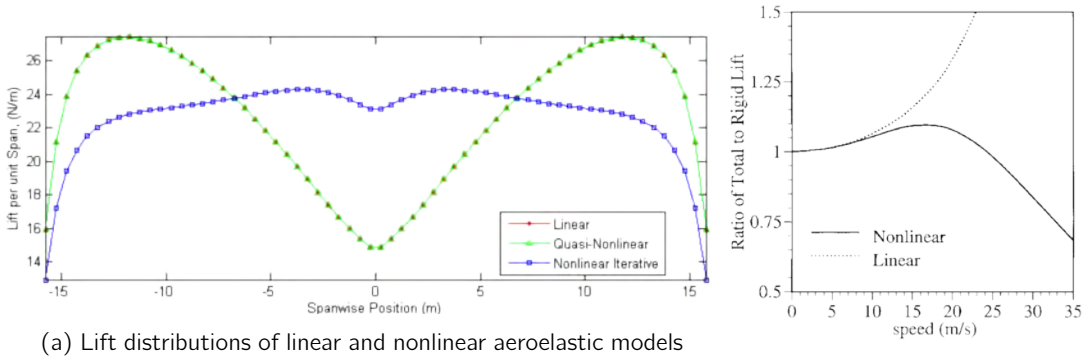

(a) Lift distributions of linear and nonlinear aeroelastic models of an identical wing, trimmed to the same total lift coefficient [15].

(b) Total lift produced by linear and nonlinear flexible wings, relative to rigid wing [75]

Figure 1.7: Examples of poor lift prediction in geometrically linear aeroelastic analysis.

The above effect demonstrates that large structural displacements necessitate the inclusion of geometric nonlinearity in the aerodynamic model to correctly capture the nature of aerodynamic forces as follower forces. In turn, these follower force effects necessitate the use of geometrically nonlinear structural models. Figure 1.8 shows the deflection of multiple nonlinear cantilever models subject to both vertical and follower tip loads.

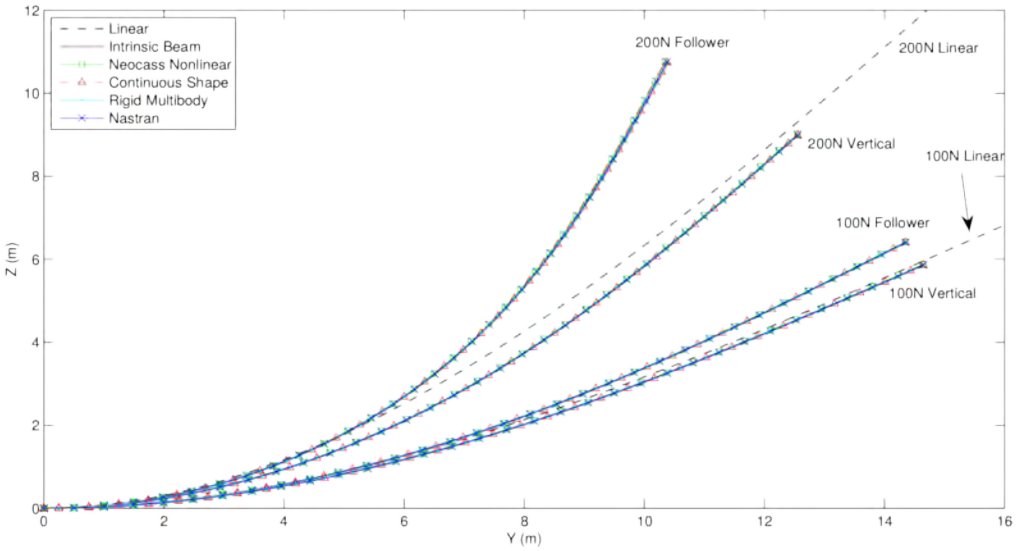

Figure 1.8: Linear and nonlinear cantilever models subject to a tip force, demonstrating the artificial extension of the linear model and the effect of follower forces on the nonlinear model [15].

The beam subject to the follower force undergoes greater displacement and rotation, resulting in a significantly shorter effective span. This is because the tip force remains perpendicular to the beams axis, thus inducing the maximum possible bending moment. Conversely, in the non-follower force case, as the beam deforms, the tip force becomes increasingly aligned with the axis of the beam, resulting in a lower bending moment in the beam and thus lower deflections. Although not shown in figure 1.8, Ritter and Cesnik [76] demonstrated that the use of a follower force with a linear FE model results in lower deformation. This is because, in the undeformed configuration, only vertical force components induce bending whilst any inboard force component acts purely along the axis of the beam, it's stiffest direction. Each node of the linear beam therefore undergoes almost purely vertical displacement and as a result there is a fictitious increase in the overall length of the beam. In aeroelastic models which compute aerodynamic forces on the deformed wing geometry, this fictitious lengthening and subsequent increase in wing area can result in unrealistically high aerodynamic forces.

Up to this point, lift has been considered the only aerodynamic force responsible for aeroelastic loading and deformation. In most cases this is a valid assumption as, in an undeformed configuration, drag forces attempt to bend the wing in its very stiff chordwise direction. However, when a wing undergoes significant vertical displacements, drag forces on the outboard portions of the wing gain a moment arm around the wing root and can result in significant torsional loads. Howcroft et al. [15] demonstrated this effect using an unsteady VLM method which employs rough drag estimates and showed that their inclusion reversed the direction of the wing root twisting moment.

Both Garcia [9] and Smith et al. [77] demonstrated this effect at higher fidelity by coupling nonlinear FE beams to CFD models of highly flexible wings. Garcia modelled an unswept wing with an aspect-ratio of 29 in transonic flow and, as shown in figure 1.9, found that a linear structural model predicted a small amount of pitch down torsion along the wing. In contrast, the nonlinear structural model showed torsional moments approximately an order of magnitude greater and in the pitch up direction, resulting in washin along the wing and increased lift which leads to a roughly 10% increase in total lift and a 25% increase in wing root bending moment.

Garcia also simulated a version of the same wing with 5° backward sweep. The results from that analysis showed that due to the inherent bend-twist coupling of a swept wing, the greater bending in the nonlinear structure results in more geometric washout. This additional washout more or less cancelled out the torsional drag effects but the nonlinear structure still shows around 10% less washout at the wingtip, leading to a noticeable increase in loading.

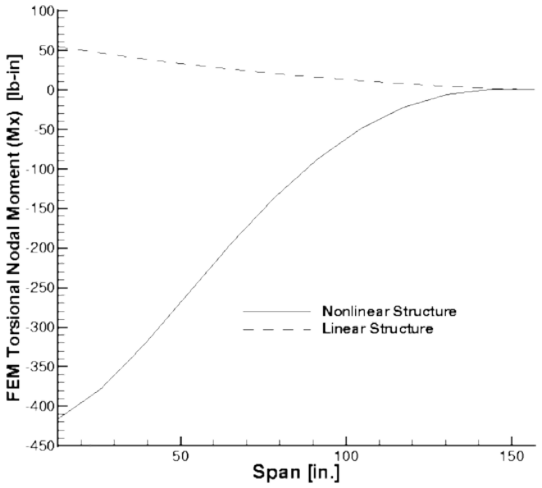

(a) Garcia [9] found that torsion due to drag resulted in an order of magnitude greater pitch up wing root moment in a nonlinear model.

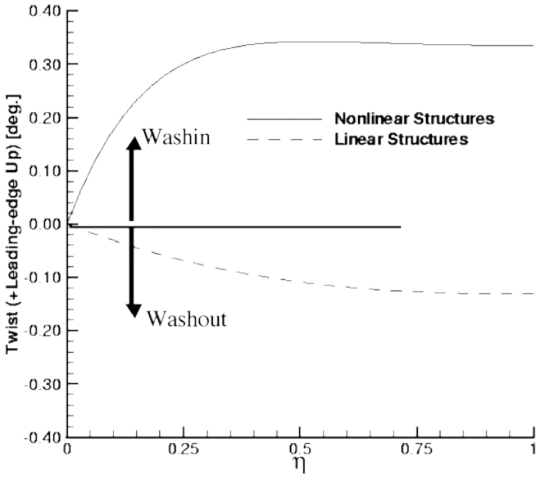

(b) As a result, the nonlinear model twists in the opposite direction and by a greater amount.

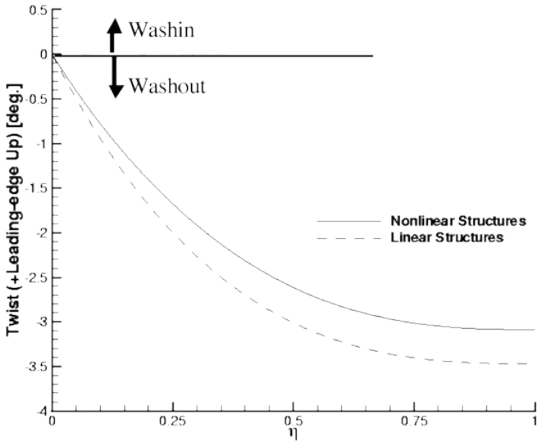

(c) In the backward swept wing, the aeroelastic twist of the was dominated by the geometric bend-twist coupling, resulting in more similar twist distributions.

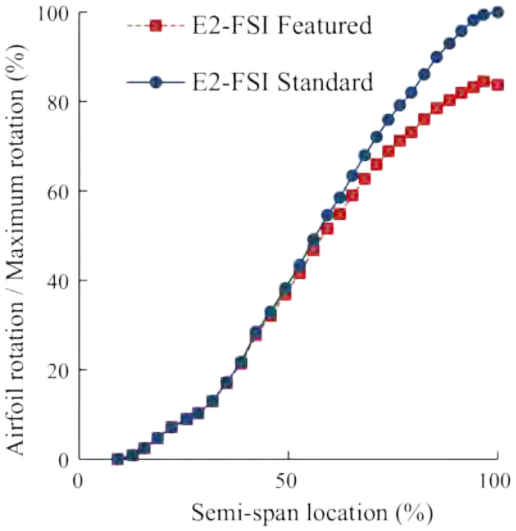

(d) Verri et al. [78] saw the same effect in their high fidelity analysis of an Embraer regional jet. (Featured = nonlinear)

Figure 1.9: Garcia [9] and Verri et al. [78] both demonstrated the important effect of drag forces on wing torsion under large displacements.

To date, there have been few published examples of aerostructural analysis using both high-fidelity CFD and high-fidelity geometrically nonlinear structural models. Medeiros et al. [20] developed a method for constructing geometrically nonlinear reduced order models of high fidelity wingbox structures which were then coupled to a RANS CFD code for both static and dynamic analysis. The work however offers little discussion of the differences between linear and nonlinear structural models. To the best of my knowledge, the recently published work of Verri et al. [78] is the only published work coupling a full-order geometrically nonlinear wingbox model to high-fidelity CFD. The work considers analysis of an Embraer regional jet with a wing aspect ratio of 12, around the upper end of current day commercial aircraft aspect-ratios, and compares aerostructural results from a 2.5g pull-up manoeuvre with both linear and nonlinear structures. Although the authors state that at this load level the wing deflection is within what would typically be considered the realm of linear behaviour, the results show some of the same geometrically nonlinear phenomena discussed to this point although the analysis provided in the paper is somewhat

limited. The geometrically nonlinear wing shows an 8% lower tip displacement and around 20% less tip washout. As a result, the aircraft in the nonlinear case is trimmed at a slightly lower angle of attack and shows a more outwardly shifted lift distribution which results in a more than 10% increase in bending and shear loads in the outboard sections of the wing. Unfortunately the work offers no analysis of stress distributions in the wingbox.

All the works mentioned to this point have considered the effects of geometric nonlinearity on aeroelastic performance. However, for these results to be of practical use we must consider at what point these nonlinearities necessitate consideration in the design process and how highly flexible wings should be designed to account for them. Although tools like PROTEUS have been used to optimise wing structures based on geometrically nonlinear aeroelastic analysis, few works have explicitly investigated the differences in wing designs optimised with and without considering such effects. Lupp and Cesnik [19] performed a basic aerostructural optimisation using the UM/NAST aeroelastic framework, optimising the structural thickness and span of a wing subject to constraints on wing root bending moment and flutter speed. The work's main focus was on the effect of including geometric nonlinearity in the flutter analysis, finding that a geometrically nonlinear flutter constraint resulted in a lower aspect-ratio wing.

A far more extensive investigation on the design implications of geometrically nonlinear behaviour was performed in two recent papers by Calderon et al. [22, 79]. Their investigations were performed using the low-fidelity aeroelastic framework NeoCASS, which couples geometrically nonlinear VLM and beam models. Using NeoCASS, Calderon et al performed a series of studies where the sizing of a wing's panels and stiffeners was optimised based on failure prediction methods typical of the conceptual aircraft design stage. The results of the work consistently showed that wings optimised using nonlinear structural analysis are lighter and longer. For a benchmark wing with an aspect-ratio of 18, the geometrically nonlinear model experienced lower loads and thus ended up 5% lighter, with a root bending moment 8% less and root torsion 50% less than the wing sized using the linear structural model, resulting in a modest 1.3% improvement in Breguet range. It should however be noted that the aerodynamic model used in this study did not include drag as a structural force and so could not have captured the drag-torsion effects described earlier in this section which greatly affected wing torsion.

Calderon et al. then performed a series of optimisations with wings ranging from aspect ratios of 10 to 26. As figure 1.10 shows, over the entire range of aspect ratios studied, the wings optimised using geometrically nonlinear analysis were lighter, achieved the same aerodynamic efficiency, and therefore had greater ranges than those optimised using the linear model. Also of interest is that the optimum aspect ratio in terms of range is higher in the nonlinear case and that only the wings at the bottom end of the aspect ratio range were subsequently found to be critically sized by gust loads. These results suggest that, although the total gain in range is modest (≈1%), geometrically nonlinear structural analysis does result in significant differences in structural sizing. When performing sizing using a full wingbox model rather than a beam model, it could be expected that these differences will be increased by the presence of stress stiffening effects like Brazier loading.

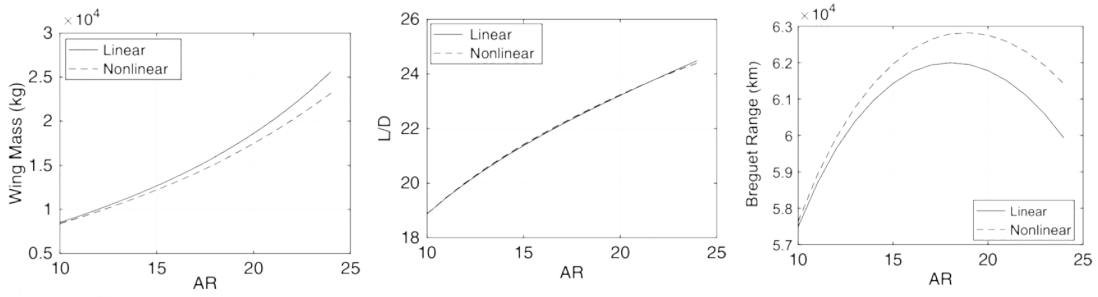

(a) Mass. (b) L/D. (c) Range.

Figure 1.10: Calderon et al. [22] found that wings optimised using geometrically nonlinear structural analysis were lighter over the entire range of aspect ratios studied, had roughly equal aerodynamic efficiency, and had a higher optimum aspect ratio than those optimised based on linear structural analysis.

1.2. Thesis Motivation and Aims

Table 1.1: A summary of the capabilities of some of the aerostructural and aeroelastic analysis and optimisation work covered in this review along with the proposed contribution of this thesis.

Author/Tool	Aerodynamic Model[1,2]	Structural Model[2,3]	Optimisation Variables		
			Structure	Planform	Shape
Howcroft et al. [15]	VLM/DLM	NL Beam			
Garcia [9], Bartels et al. [14]	RANS	NL Beam			
Verri et al. [78]	RANS	NL Shell			
Calderon et al. [22, 79]	VLM	NL Beam	✓		
OpenAeroStruct [55, 56, 58]	VLM	Beam	✓	✓	
FEMWET [57]	Q3D	Beam	✓	✓	✓
PROTEUS [21, 23, 60, 63, 68–70]	VLM	NL Beam	✓		
Jovanov [72]	Euler	NL Beam	✓		
Stanford et al. [61, 64, 65, 67, 74, 80]	VLM	Shell	✓		
Stanford [66]	VLM	Shell	✓	✓	✓
UM/NAST [19]	ST	NL Beam	✓	✓	
Dillinger et al. [71]	Euler	Shell	✓		
MACH [2–6]	RANS	Shell	✓	✓	✓
This thesis	RANS	NL Shell	✓	✓	✓

[1] Aerodynamic models: DLM–Doublet lattice method, VLM–Vortex lattice method, ST–Strip theory, Q3D–Quasi 3D, Euler–Euler CFD, RANS–RANS CFD
[2] Fidelity: Low, Mid, High
[3] Structural models: NL–Nonlinear

Table 1.1 summarises the capabilities of the aerostructural analysis and optimisation tools just described and demonstrates the primary motivation for this thesis. There is currently a knowledge gap in the detailed aerostructural design of high-aspect-ratio wings (HARW) resulting from a capability gap between low-fidelity aeroelastic optimisation frameworks that include geometric nonlinearities and high-fidelity aeroelastic optimisation frameworks that cannot.

The work done to date on the effect of geometric nonlinearity on the structural sizing of HARW has either involved structure-only optimisation using low-fidelity tools [22] or single analyses with high-fidelity tools which are currently too computationally expensive and lack the gradient calculation capability required for optimisation. [78]. As such, it is not clear what the effects of geometric nonlinearity are on optimal design of aircraft wings and at what point they become an important consideration.

The aim of this thesis is therefore twofold:

1. To develop a tool capable of performing high-fidelity aerostructural analysis and optimisation using RANS CFD and a geometrically nonlinear full wingbox FE models.
2. To use the tool to investigate the consequences of geometrically nonlinear effects on the aerostructural design of modern transport aircraft.

Specifically, this thesis work has been performed at the Univeristy of Michigan's MDO Lab and thus the development portion of this thesis focuses on expanding the MACH framework's capabilities to include a

geometrically nonlinear finite element formulation.

To tackle the second aim, the second portion of this thesis involves a series of first structural, then aerostructural, analyses and optimisation studies of the undeflected common research model (uCRM)s, a pair of benchmark aerostructural models developed by Brooks et al. [81], representative of both moderate and high aspect-ratio commercial transport aircraft. These studies will allow for the identification of differences due to purely structural nonlinearities under fixed loading and their consequences on optimal designs, before the aerostructural studies assess the consequences of combined structural and aeroelastic nonlinearities. Of particular interest is whether these nonlinearities result in optimised wings with more or less passive load alleviation as the work presented so far from Lupp and Cesnik [19], Garcia [9], Verri et al. [78] and Calderon et al. [22] offer contradictory results.

These aims are combined in the top level research question:

> Is it computationally feasible to use geometrically nonlinear finite element formulations in high fidelity aerostructural design optimisation? And, if so, do geometric nonlinearities have a significant effect on the optimal aerostructural design of high aspect ratio wings for future commercial aircraft?

This can be broken down into the following sub-questions:

1. What is the computational cost of performing high-fidelity aerostructural analysis and optimisation with a geometrically nonlinear structural model and how can this cost be minimised?
   (a) How can the robustness and efficiency of a basic Newton-Raphson solver be improved for geometrically nonlinear analysis of high aspect-ratio wingboxes?
      i. Which line search strategy provides the required solver robustness for the minimum additional cost? Is this choice problem dependent?
      ii. Do approximate Jacobian update strategies reduce the solver's overall computational cost without sacrificing robustness?
      iii. Are load incrementation or arc-length control methods necessary for modelling wingboxes not undergoing buckling? If so, how can incrementation be controlled to minimise computational cost?
   (b) How can MACH's aerostructural solver be modified to work with the TACS nonlinear solver?
      i. If required for robustness in the structural solver, how can load incrementation be incorporated into the aerostructural solver?
   (c) What is the overall cost penalty associated with the inclusion of structural geometric nonlinearity in structural and aerostructural analysis and optimisation?
      i. Does the increased cost make aerostructural optimisation with geometrically nonlinear structural analysis impractical?
      ii. Where might the most performance improvements be possible in the future?
2. What effects do structural geometric nonlinearities have on the results of structural and aerostructural analysis and optimisation of HARW?
   (a) How does geometric nonlinearity affect the results of structural and aerostructural analysis?
      i. How does geometrically nonlinearity affect the deformations and stresses in a wingbox for a given loading?
      ii. How does geometric nonlinearity affect wing deformation during flight?
      iii. How does this affect the wing's passive load alleviation potential and consequently the load and stress distribution in the wingbox?
      iv. How does this affect the overall lift, drag and trim of an aircraft?

v. Which of these effects are due to purely structural nonlinearities and which are due to aeroelastic nonlinearities?

(b) How does geometric nonlinearity affect the results of structural and aerostructural optimisation?

i. How do these effects alter the optimal sizing of a wingbox for a given loading?

ii. How do these effects alter the optimal wing design in terms of structural sizing, wing planform and wing shape?

iii. Which of these design changes are due to purely structural nonlinearities and which are due to aeroelastic nonlinearities?

(c) How critical is the inclusion of geometric nonlinearity in the optimal design of HARW for future commercial transport aircraft?

i. At what point do these effects become important? Is aspect-ratio the main determining factor? Or are these effects also present in low to moderate aspect ratio wings?

ii. Do the differences in optimised designs make aerostructural optimisation with linear finite element models invalid for the conceptual, preliminary or detailed design of HARW for future commercial transport aircraft?

1.3. Thesis Overview

The remainder of this thesis is structured as follows.

Chapter 2 introduces the various tools within MACH used throughout this work.

Chapter 3 gives some background on solution techniques for nonlinear FE and fluid-structure interaction (FSI) problems before describing the development, validation and testing of the nonlinear solver developed in this work for MACH's FE code, TACS and the modifications made to the framework's aerostructural solver.

Chapter 4 describes both the setup of the uCRM models and the comparisons of structural and aerostructural analysis performed on them with linear and nonlinear analysis.

Chapter 5 moves onto to a series structural, and finally aerostructural, optimisations and seeks to link the differences in optimal designs back to the nonlinear effects demonstrated in chapter 4

Chapter 6 restates the findings from each chapter, draws some overall conclusions in order to answer the above research questions and makes recommendations for future investigations.

2

The MACH Framework

In the past decade, the MDO Lab has developed the MACH (MDO of Aircraft Configurations at High-fidelity) framework, which represents the current state of the art in high-fidelity aerostructural optimisation. The framework includes CFD and FE solvers with efficient adjoint derivative implementations along with fully differentiated modules required for coupled analysis and optimisation (e.g. geometry parameterisation, load and displacement transfer, and mesh warping) which are used at various points throughout this work. Figure 2.1 shows the extended design structure matrix (XDSM) for a typical aerostructural optimisation problem as will be performed in this work. In this chapter I briefly review the tools used in all stages of this aerostructural analysis and optimisation chain, where useful for later understanding, I cover the technical details of some of the tools in more depth.

2.1. Aerodynamic Analysis: ADflow

The flow solver in MACH is ADflow[1], a finite-volume CFD solver for structured multiblock and overset meshes [82]. ADflow solves the compressible Euler, laminar Navier-Stokes, and RANS equations with a second-order accurate spatial discretisation. The solver employs a variety of numerical methods to converge to a steady-state solution, including multigrid, approximate Newton-Krylov, and Newton-Krylov algorithms [83]. The combination of these various iterative methods makes ADflow robust and fast. ADflow also solves the discrete adjoint equations, enabling efficient computation of derivatives independent of the number of design variables. The solution of the discrete adjoint in ADflow relies on the ADjoint approach, which uses algorithmic differentiation (AD) to compute partial derivatives and a Krylov method to solve the linear system [84].

[1] `github.com/mdolab/adflow`

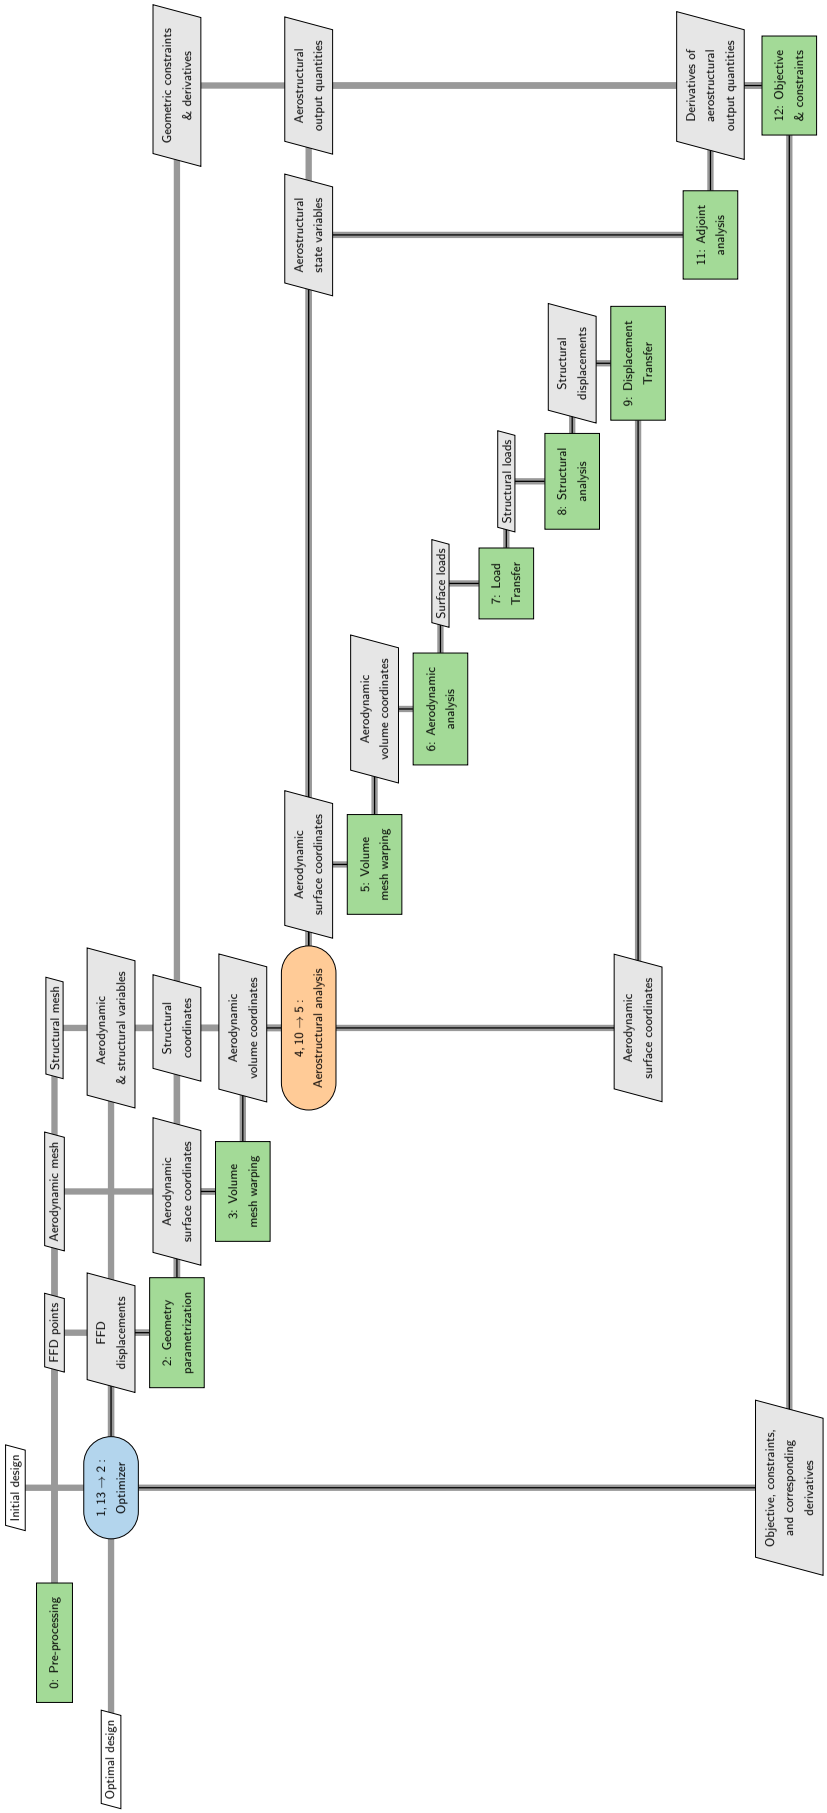

Figure 2.1: The XDSM for a typical aerostructural optimisation problem in MACH

2.2. Structural Analysis: TACS

The structural solver used in MACH, and throughout this work, is TACS (toolkit for analysis of composite structures)[2] [42]. TACS was developed specifically for the kinds of poorly conditioned FE problems typical of thin shell structures like wingboxes, which may have condition numbers of $\mathcal{O}\left(> 10^9\right)$ [85]. This poor conditioning makes the linear systems representing shell structures difficult or impossible to solve for the Krylov based iterative solvers typically used for large linear systems. To get around this problem, TACS forms an essentially perfect preconditioner for a Krylov solver using a direct solution method. For linear structural analysis, the expensive factorisation of the preconditioner can be performed once for a given structure and then reused to solve for many different load cases with very little computational cost. This is a particularly useful characteristic during coupled aerostructural analysis with a linear structure [6].

TACS also computes gradients using the adjoint method but, unlike ADflow, all partial derivative computations are analytic.

2.2.1. Constraint Aggregation

In a structural or aerostructural optimisation problem, constraints are almost always imposed on the maximum allowable stress in a structure. Adding an individual constraint for each point in the structure at which stress is computed would lead to $\mathcal{O}\left(10^4 - 10^7\right)$ constraints, making optimisation impractical, while simply constraining the maximum stress value leads to poor optimiser behaviour as the derivative of the max function is discontinuous [86]. To address this issue, TACS offers a selection of constraint aggregation techniques in Kreiselmeier-Steinhauser (KS) and P-norm aggregation. All of which aim to provide a smooth, differentiable approximation of the maximum value of a set of functions. Although Lambe and Martins [87] recommend the use of the P-norm aggregate, I use KS aggregation throughout this work as it is the only method implemented in TACS which uses average element stresses rather than the values computed at integration points. I will show later, in chapter 5, that this is critical for reducing the adverse effect of stress singularities on optimisation problems.

2.2.2. The bladeFSDT Constitutive Model

As described in chapter 1, buckling is the onset of an instability in a structure and is caused by compressive stress stiffening effects. Due to their strong stress stiffening behaviour, buckling is a critical failure mode in thin shell structures such as wingboxes [88, 89]. As a result, the panels in wingboxes are reinforced with stiffeners (or stringers) whose sizing is almost entirely dictated by buckling behaviour. In TACS, stringers can either be explicitly modelled using shell or beam elements or, more commonly, using a constitutive model referred to as the bladeFSDT model. As the name suggests, the bladeFSDT model is based on first-order shear deformation theory, which assumes that stain components vary linearly through the thickness of the shell [85]. The model includes the effect of blade stiffeners by adding equivalent stiffness terms to the skin shell elements, essentially 'smearing' the stiffness of stiffeners over the entire panel area. As shown in figure 2.2, the bladeFSDT model is parameterised by 3 additional design variables per panel for the thickness, height and pitch of the stiffeners in addition to the original skin thickness variable. This parameterisation is one of the primary advantages of the bladeFSDT approach over explicitly modelled stiffeners, which can only easily be given thickness design variables.

Failure criteria are analysed at 3 locations through the thickness of the virtual panel. The upper and lower surface of the skin, which are in a 2D plane stress state, and at the extrema of the stiffener blade, which is assumed to be under purely axial stress. The failure criteria evaluated depends on the material being modelled, with Von Mises stress used for isotropic metallic structures, and either a maximum strain or Tsai-Wu criterion for composites.

Another advantage of the bladeFSDT model over explicit stiffeners is the ability to predict buckling failure without the need for any additional eigenvalue based analyses. To do this, TACS defines a force based

[2] `github.com/gjkennedy/tacs`

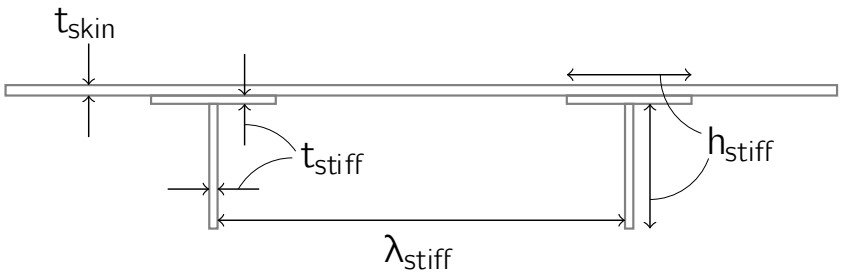

Figure 2.2: The parameterisation of TACS' bladeFSDT constitutive model.

failure envelope:

$$\frac{F_x}{F_{x,cr}} + \frac{F_{xy}^2}{F_{xy,cr}^2} \leq 1 \tag{2.1}$$

Based on critical compressive, $F_{x,cr}$, and shear, $F_{xy,cr}$, loads in the local axis of the panel stiffeners. At each element these critical loads are calculated using the empirical relations presented by Stroud and Agranoff [90] for global panel buckling, inter-stringer skin buckling and stringer buckling of an infinitely wide flat panel of the same length as the wingbox panel the element is a part of.

2.2.3. Python Interface: pyTACS

TACS is written in C++ but interfaces with the MACH framework through a Python interface known as pyTACS. pyTACS contains all the classes and methods necessary for the definition and solution of a structural problem, including performing common linear algebra operations on TACS' distributed vector and matrix objects. As such the solution process can be controlled in great detail at the Python level, whilst the intensive numerical work is still handled in a highly efficient and parallelised manner in the C++ layer.

2.3. Load and Displacement Transfer: Rigid Link

Since the meshes used in the aerodynamic and structural analyses typically do not match each other and are often even non-coincident, a method is required both for transferring structural displacements to the aerodynamic surface and for transferring aerodynamic forces to the structural mesh. The load-displacement transfer method used in MACH is known as the rigid link transfer (RLT) and was first proposed by Brown [91]. Figure 2.3 demonstrates the RLT method for displacement transfer. The method works by first finding the nearest point on the structural mesh for each aerodynamic surface node. The translation and rotation of the base of the link is interpolated using TACS' own shape function before the displacement of the aerodynamic node is then computed based on the rigid translation and rotation of the link.

Using this method, the aerodynamic nodal displacements can be written as a linear function of the structural states:

$$\left\{X_S\right\} = \left\{X_J\right\} + \left\{u_A\right\} \tag{2.2}$$

$$\left\{u_A\right\} = \left\{u_t\right\} + \left\{u_r\right\} \times \left\{r\right\} \tag{2.3}$$

$$\left\{u_A\right\} = \left[T\right] \left\{u\right\} \tag{2.4}$$

By conserving the work done on both meshes, it can then be shown that the transfer of forces must be performed using the transpose of the displacement transfer:

$$\left\{F\right\} = \left[T\right]^T \left\{F_A\right\} \tag{2.5}$$

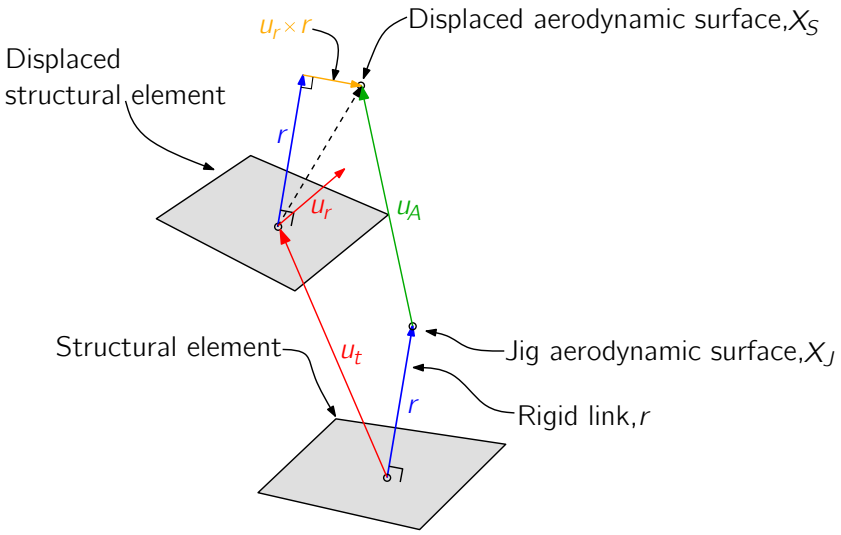

Figure 2.3: The RLT method for transferring displacements from the structural to aerodynamic mesh [6].

It should be noted that this method is not accurate for large rotations as the term computing the aerodynamic displacement due to structural rotations, $\left\{u_r\right\} \times \left\{r\right\}$, is a linearised approximation. For large rotations this term should be replaced by a rotation matrix, R, giving:

$$\left\{u_A\right\} = \left\{u_t\right\} + \left[R\left(u_r\right)\right]\left\{r\right\} \tag{2.6}$$

In this work, however, I retain the linearised form of the displacement transfer scheme as modifying the source code for the load and displacement transfers and their derivatives was not possible within the time available. This simplification is not ideal and results in a slight fictitious increase in the chord length of wing sections under twisting but I consider it acceptable as, for the most part, the structural and aerodynamic meshes are coincident, rendering the rotational term in equation 2.3 negligible.

The RLT method is implemented within the TACS source code and so, like TACS, computes gradients analytically.

2.4. Mesh Warping: IDWarp

During aerostructural analysis and optimisation, the aircraft's outer mould line (OML) deforms due to both design changes and structural deformations. As these surface deformations occur, the CFD volume mesh must also be smoothly deformed to maintain high quality cells. This deformation is performed in MACH by a tool called IDWarp[3] [92]. IDWarp uses an inverse-distance weighting method as proposed by Luke et al. [93]. The advantage of this method is that it maintains good orthogonality in near-wall cells, which is vital for RANS meshes, at a significantly lower computational cost than elasticity based approaches. This cost is vital for efficient aerostructural analyses as the mesh must be deformed every time structural displacements are transferred to the CFD mesh. Derivatives of the volume mesh coordinates with respect to surface coordinates are computed using reverse-mode AD.

2.5. Coupled Aerostructural Analysis: pyAeroStructure

PyAeroStructure interacts with the aforementioned tools in order to coordinate coupled aerostructural analyses and gradient computations. The module can perform both the analysis and gradient computation with either block partitioned Gauss-Seidel or monolithic Newton-Krylov approaches. As is explained

[3] `github.com/mdolab/idwarp`

further by Kenway et al. [6], using a matrix free Newton-Krylov method means the monolithic solver never explicitly forms the full coupled system of equations and can instead reuse the residual evaluation and preconditioning procedures from each individual solver, massively reducing the implementation cost. Similarly, the monolithic coupled adjoint solver reuses the adjoint solvers of TACS and ADflow.

2.6. Geometric Parameterisation and Constraints: pyGeo

Whilst ADflow and TACS each control a set of design variables that affect only their own models, the geometric parameterisation of a given aircraft affects the geometry of both the aerodynamic and structural models and is handled by a separate module, pyGeo[4]. pyGeo [43] uses a free-form deformation (FFD) [94] approach for geometric parameterisation. The FFD approach embeds a baseline geometry to be parameterised in a series of control volumes defined by a set of control points, deformations can then mapped from the control points to the surface mesh of the geometry using any number of interpolation or 'mapping' schemes. Kenway [95], who implemented the FFD scheme in pyGeo describes some of its key advantages as the consistency of it's parameterisation across disciplines and the ability to efficiently calculate derivatives of deformations with respect to control point displacements. Components of control point displacements can be controlled individually to offer local shape control and also grouped together to allow single design variables to control global wing shape parameters like span and sweep. Figure 2.4 demonstrates this global-local geometry parameterisation approach, showing the effect of FFD design variables which control wing span, sweep, section twist and local shape.

Derivatives in pyGeo are computed using a mixture of analytic and complex-step methods.

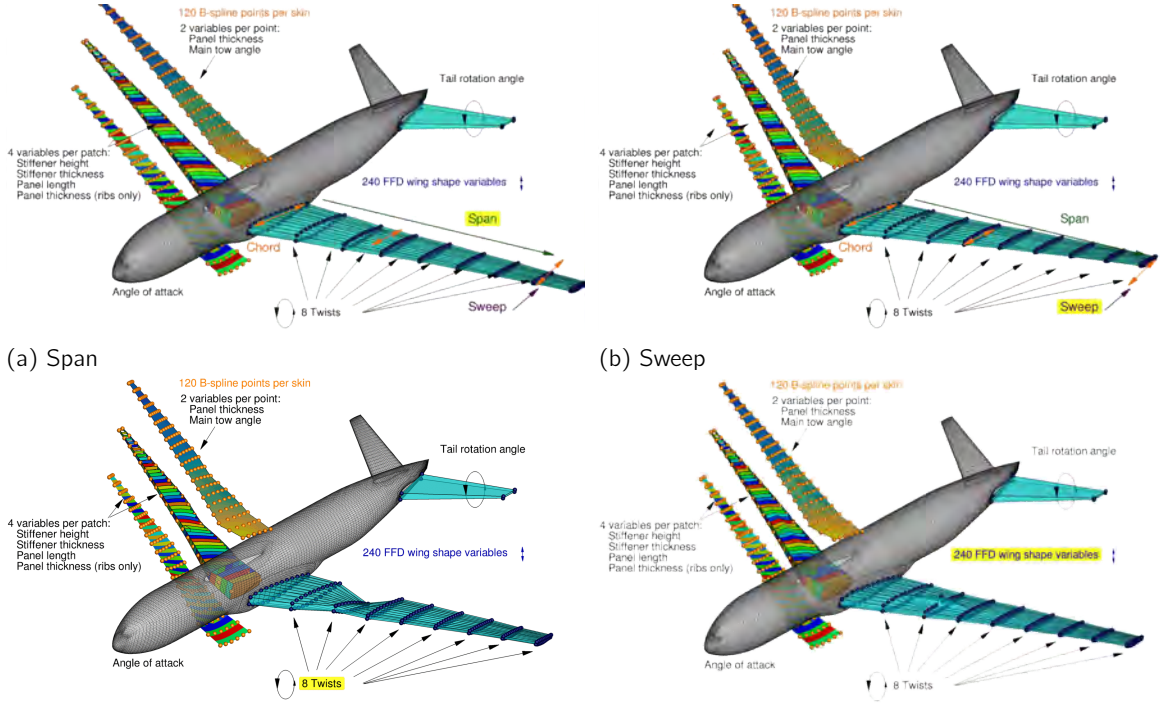

(a) Span

(b) Sweep

(c) Twist

(d) Local

Figure 2.4: Examples of both global and local geometric design variables typically used in MACH [2]

[4] github.com/mdolab/pygeo

3

Solver Development

In this chapter I describe the new solvers developed to enable geometrically nonlinear structural and aerostructural analysis in MACH. Section 3.1 describes the new solver developed for nonlinear analysis in TACS while section 3.2 covers the modifications made to MACH's coupled aerostructural solver. I each section I provide a small amount of background on relevant nonlinear solution methods before describing the implemented solvers and presenting the results of some performance tests, comparing solution times between different nonlinear methods and between linear and nonlinear analyses. Finally, in section 3.4 I summarise the key takeaways from the development of both solvers.

3.1. Structural Solver

3.1.1. Nonlinear Finite Element Equations

In any structural finite element (FE) problem, we seek to solve a system of residual equations which driving the imbalance between internal and external forces to zero:

$$\left\{\mathcal{R}\left(u\right)\right\} = \left\{F_{in}\left(u\right)\right\} + \left\{F_{ex}\left(u\right)\right\} = 0 \tag{3.1}$$

In a linear structural analysis, the equilibrium of forces is computed in the undeformed configuration of the structure only. External forces, $\left\{F_{ex}\right\}$, are assumed to be constant and, using small strain assumptions, the internal forces, $\left\{F_{ex}\right\}$, can be written as a linear function of the structural displacement states $\left\{u\right\}$:

$$\left[K\right]\left\{u\right\} = \left\{F_{ex}\right\} \tag{3.2}$$

This linear system can be solved extremely efficiently by modern computers and, as is done in TACS, the expensive factorisation of $\left[K\right]$ can be performed once and the factored version of the matrix reused to solve for multiple loadings [42]. This is a particularly useful strategy during aerostructural analysis where the aerodynamic loads are continually updated based on the deformed wing shape.

For a geometrically nonlinear FE problem we must however stick with equation 3.1, where both the internal and external forces are nonlinear functions of the displacement. The dependence of the external forces on the displacements $\left\{u\right\}$ is due to follower force effects, while the nonlinearity of the internal forces comes from the Green-Lagrange strain tensor:

$$\epsilon_{ij} = \frac{1}{2}\left(\frac{\partial U_i}{\partial X_j} + \frac{\partial U_j}{\partial X_i} + \frac{\partial U_k}{\partial X_i}\frac{\partial U_k}{\partial X_j}\right) \tag{3.3}$$

Where U is the displacement vector ($U = \left\{u, v, w\right\}^{\mathrm{T}}$ in 3D) and X is the coordinate vector ($X = \left\{x, y, z\right\}^{\mathrm{T}}$ in 3D).

To demonstrate the utility of the Greene-Lagrange strain measure consider the case shown in figure 3.1, where a simple bar, originally aligned with the x-axis undergoes pure vertical displacement at one end. Intuitively, since the length of the deformed bar is greater than the undeformed bar, we expect to measure a positive axial strain. As shown in the figure, the deformation U of an arbitrary point on the bar which started at point $(x_0, 0)$ is $(u, v) = (0, \Delta y x_0/L)$. The axial strain at this point is given by the first component of the Green-Lagrange strain tensor:

$$\epsilon_{xx} = \frac{1}{2}\left(\frac{\partial u}{\partial x} + \frac{\partial u}{\partial x} + \frac{\partial u}{\partial x}\frac{\partial u}{\partial x} + \frac{\partial v}{\partial x}\frac{\partial v}{\partial x}\right) \tag{3.4}$$

$$\epsilon_{xx} = \frac{\partial u}{\partial x} + \frac{1}{2}\left(\left(\frac{\partial u}{\partial x}\right)^2 + \left(\frac{\partial v}{\partial x}\right)^2\right) \tag{3.5}$$

$$\epsilon_{xx} = 0 + \frac{1}{2}\left(0 + \left(\frac{\Delta y}{L}\right)^2\right) = \frac{1}{2}\left(\frac{\Delta y}{L}\right)^2 \tag{3.6}$$

Equation 3.6 shows that the Green-Lagrange strain is positive as expected however, the linearised strain measure (the $\partial u/\partial x$ term in equation 3.5) incorrectly predicts zero strain. Similarly, for a bar undergoing a rigid rotation, it can be shown that the linearised strain measure predicts a fictitious compressive strain whereas the nonlinear Green-Lagrange strain measure correctly gives zero strain.

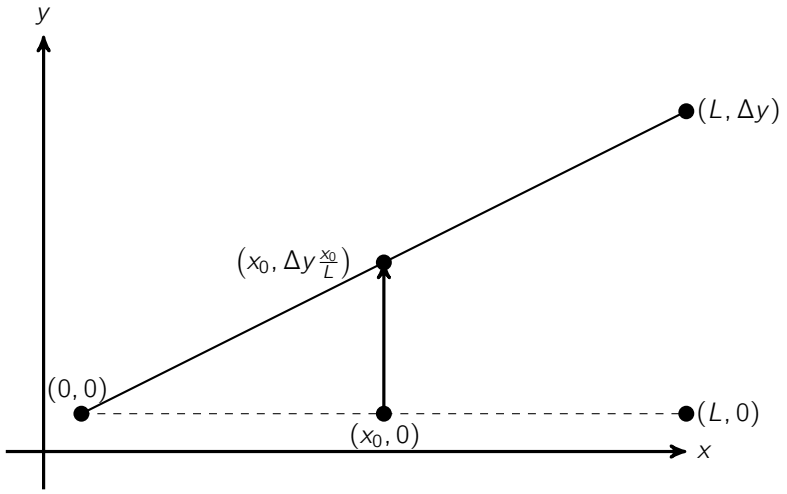

Figure 3.1: Bar element undergoing pure vertical displacement

Additionally, for elements with rotational degrees of freedom, such as beam and shell elements, the kinematic relationships which define the contribution of the rotational degrees of freedom to interpolated displacements inside each element must be accurate for finite rotations. In linear analysis these kinematic relationships are linearised based on small angle assumptions and are thus invalid for moderate or large rotations.

TACS contains two geometrically nonlinear shell element formulations. A 'largerot' formulation that is fully geometrically nonlinear and a 'nonlinear' formulation that uses the Green-Lagrange strain relationship but retains the linear kinematic relationships. The 'nonlinear' formulation is therefore only valid under small rotations but still captures stress stiffening effects correctly so would be able to replicate the results in figures 1.5 and 1.6 but not those shown in figure 1.4. Since the aim of this thesis is to accurately model HARW under large displacements, I use the 'largerot' formulation throughout and any reference to 'nonlinear' analysis refers to the use of this formulation.

3.1.2. Base Newton algorithm

The implemented solver, like almost all algorithms for solving nonlinear FE problems, is based on the Newton-Raphson method. The Newton-Raphson method is an iterative approach that involves the linearisation of the residual equations around the current point:

$$\left\{\mathcal{R}\left(u_{i+1}\right)\right\} \approx \left\{\mathcal{R}\left(u_i\right)\right\} + \left[K_T(u_i)\right]\left\{\Delta u_i\right\} \tag{3.7}$$

Where $u_{i+1} = u_i + \Delta u_i$ and the tangent stiffness matrix, K_T, is the Jacobian of the residual equations with respect to the nodal displacement state variables:

$$\left[K_T(u_i)\right] = \left[\frac{\partial \mathcal{R}}{\partial u}\right]|_{u=u_i} = \left[\frac{\partial F_{in}(u)}{\partial u}\right]\Big|_{u=u_i} + \left[\frac{\partial F_{ex}(u)}{\partial u}\right]\Big|_{u=u_i} \tag{3.8}$$

Equation 3.7 is then solved for the displacement update, Δu, that drives the linearised residuals to zero:

$$\left\{\Delta u_i\right\} = \left[K_T(u_i)\right]^{-1}\left\{-\mathcal{R}\left(u_i\right)\right\} \tag{3.9}$$

The process above is then repeated, linearising about the new point, u_{i+1}, until the residuals are below an acceptable tolerance.

Although popular, the Newton-Raphson method is not without its problems, on both robustness and efficiency grounds:

Non-convergence: The Newton-Raphson method converges reliably at a quadratic rate when in the 'basin of attraction' of a solution. However, because the calculated displacement step Δu is based on a linearisation of the residual equations, the method tends to converge very slowly or even diverge when far away from the solution or in the presence of strong nonlinearities.

Computational Expense: Every iteration of the Newton-Raphson method requires the tangent stiffness matrix to be recalculated and then factored, making the process very computationally expensive.

In an attempt to mitigate these issues, I implement a series of common improvements for nonlinear FE solvers.

3.1.3. Load incrementation

For strongly nonlinear problems, applying loads or displacements in increments can improve convergence. Incrementation is effective because it can ensure that the current state is always close enough to the solution of the current increment to guarantee reliable Newton convergence. Arc-length methods are an alternative to pure load or displacement implementation where both the displacement and load step are treated as unknowns that must be solved for [96]. Doing so allows these methods to stably traverse unstable sections of a structure's equilibrium path and are thus necessary for simulating the kind of snap-through and snap-back behaviour seen in post-buckling analyses. In this work I am concerned only with modelling pre-buckling structural behaviour and thus I use pure load incrementation. The conservative nature of the buckling constraints described in chapter 2 ensure that designs produced during optimisation remain stiff enough that the FE model itself does not explicitly buckle, which would cause non-convergence issues. I implement load incrementation by redefining the residual to include a load scaling factor λ:

$$\left\{\mathcal{R}\right\} = \left\{F_{in}\right\} + \lambda\left\{F_{ex}\right\} \tag{3.10}$$

To control the size of each load increment, I use the adaptive load stepping method of Beluni and Chulya [97], where the current load step is increased or decreased based on the actual and desired number of iterations taken to solve previous increment.

$$\lambda_i - \lambda_{i-1} = \Delta\lambda_i = \sqrt{\frac{N_{des}}{N_{i-1}}}\Delta\lambda_{i-1} \tag{3.11}$$

The desired number of iterations, N_{des}, can then be used by the user to control how bold the solver should be with load incrementation.

3.1.4. Convergence criteria/exit conditions

An increment is judged to have reached convergence when either an absolute or relative measure of the residual norm reaches a specified tolerance. The relative residual norm is normalised by the norm of the scaled external force vector.

$$\|\mathcal{R}\| < \delta_{abs} \tag{3.12}$$

$$\|\mathcal{R}_{rel}\| = \frac{\|\mathcal{R}\|}{\lambda \|F_{ex}\|} < \delta_{rel} \tag{3.13}$$

I use looser convergence tolerances for intermediate load increments in order to accelerate overall convergence.

Solver is judged to have finished when above convergence criteria are satisfied and solver is at specified maximum load scale. If the solution diverges or fails to converge an increment in a specified number of iterations, the solution is reset to last converged increment, the load increment halved, and then reapplied. Conversely, if the solver has not finished after a specified number of increments or fails to converge an increment at the minimum allowable increment size the solution is judged to have failed, in which the solver returns a failure flag which can be passed to an optimiser.

3.1.5. Line Search Methods

Applying the full displacement step in each Newton-Raphson iteration can lead to little or no convergence, especially in the early stages of a solution process. To improve this behaviour, the computed step can instead be used as a direction along which to search for a point which improves the current solution [98]. There are multiple options for both the metric by which to judge the solution improvement and the line searching algorithm. I implement 3 kinds of line search:

Minimum residual search

This method employs SciPy's 1D minimisation routine, *minimize_scalar*, to find the point along the search direction that minimises the residual norm. Using this pre-existing module greatly decreases implementation effort and allows for effortless switching in and out of the different line search algorithms provided by SciPy. In addition to the line search algorithm, the user can specify the maximum number of iterations, tolerance and upper and lower bounds on the step size.

Backtracking resdidual search

Calculating the residual still has an appreciable computational cost and so it is typically undesirable to use an exact line search technique which very accurately minimises the residual measure. Additionally it can often be beneficial to accept a slightly higher residual in return for taking a larger displacement step which moves the structure closer to the final solution. For these reasons, it is common for nonlinear solvers to use an inexact line search technique which aims merely to find a step length which reduces the residual measure 'enough'. A popular method is a backtracking search. This method works by backtracking from the full step length, until the residual norm satisfies the sufficient decrease, or Armijo-Goldstein [99], condition. This condition stipulates that the merit function, f, has decreased by at least a specified fraction μ of the amount predicted by the gradient of the function in the search direction, relative to a reference value C_i:

$$f(u_i + \alpha \Delta u_i) \leq C_i + \mu \alpha \Delta u_i^{\mathsf{T}} \nabla f \tag{3.14}$$

Where, in this case, f is the residual norm and the reference value C_i is $f(u_i)$. This formulation guarantees that each step taken by the solver reduces the residual, provided there is no lower limit on step size.

However, always requiring a decrease in the residual can lead to extremely slow convergence. To combat this I implement a non-monotone backtracking line search. In a non-monotone line search, the reference value C_i is chosen such that the merit function is allowed to increase within reason from one iteration to the next. This enables the solver to take some larger steps but not so large that the solver diverges as it would without a line search, the value of C_i should also gradually decrease with the merit function. There are two popular methods for choosing C_i:

- Grippo et al. [100] use the maximum value of f from the previous m iterations.
- Zhang and Hager [101] use a weighted average of all previous values of f, with recent terms being more heavily weighted.

These techniques were invented for optimisation algorithms where the merit function is likely to remain the same order of magnitude throughout. However, in this case, the merit function is the residual norm which may change by >6 orders of magnitude over the solution of a given increment. This means that the average residual norm value will always be dominated by the earlier values. To combat this problem I also implement a version of Zhang & Hager's method which uses a geometric weighted mean instead of an arithmetic mean.

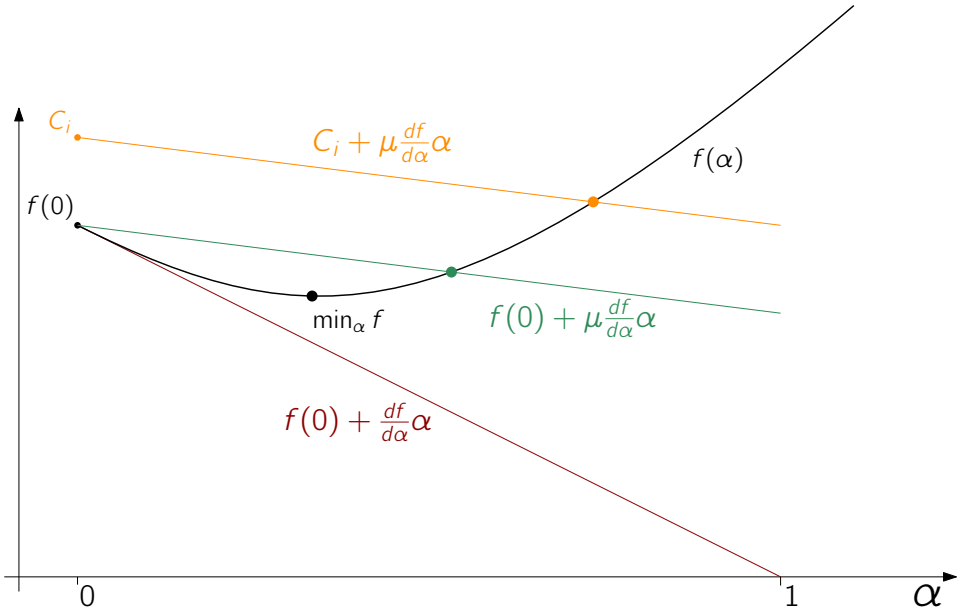

Figure 3.2: A graphical comparison of the termination criteria for the minimising, backtracking and non-monotone line searches. The black, green and orange points show the steps taken by the minimising, backtracking and non-monotone searches respectively.

Minimum energy search

Matthies and Strang [102] proposed using a line search to find the point at which the residual is orthogonal to the displacement step. In other words, finding the step size α that satisfies:

$$E = \left\{\Delta u_i\right\}^\mathsf{T} \left\{\mathcal{R}\left(u_i + \alpha \Delta u_i\right)\right\} = 0 \tag{3.15}$$

I refer to this method as the minimum energy method as the product of a force and a displacement represents a measure of energy, even if this energy does not necessarily have an intuitive physical interpretation. This minimum energy line search is implemented using SciPy's *root_scalar* root finding routine. Like *minimize_scalar*, *root_scalar* contains a variety of search methods, a simple secant method was found to perform well.

Line Search Initiation criteria

When initially testing the solver, the addition of line searches lead to robust but extremely slow convergence. An example of this for a simple tip loaded cantilever problem is shown in figure 3.3. Figure 3.3a provides an insight into this behaviour, showing the merit functions for both the residual and energy based line search methods in the first iteration of the initial load increment, where α is the step length as a fraction of the computed Newton-Raphson step. In this case, the minimum residual occurs below 1% of the original step length whilst the point of zero energy occurs around a step of 3%. Both line search methods therefore severely limit the size of steps taken by the solver, as shown in figures 3.3b and 3.3c, where an analysis using the minimum energy line search takes 108 iterations to converge.

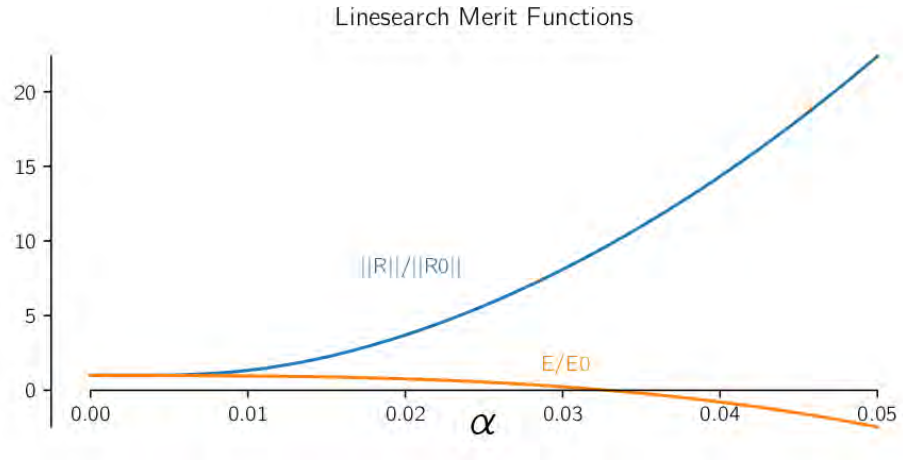

(a) In the first iteration of an increment, the optimal points on both the energy and residual line search merit functions result in very small step sizes.

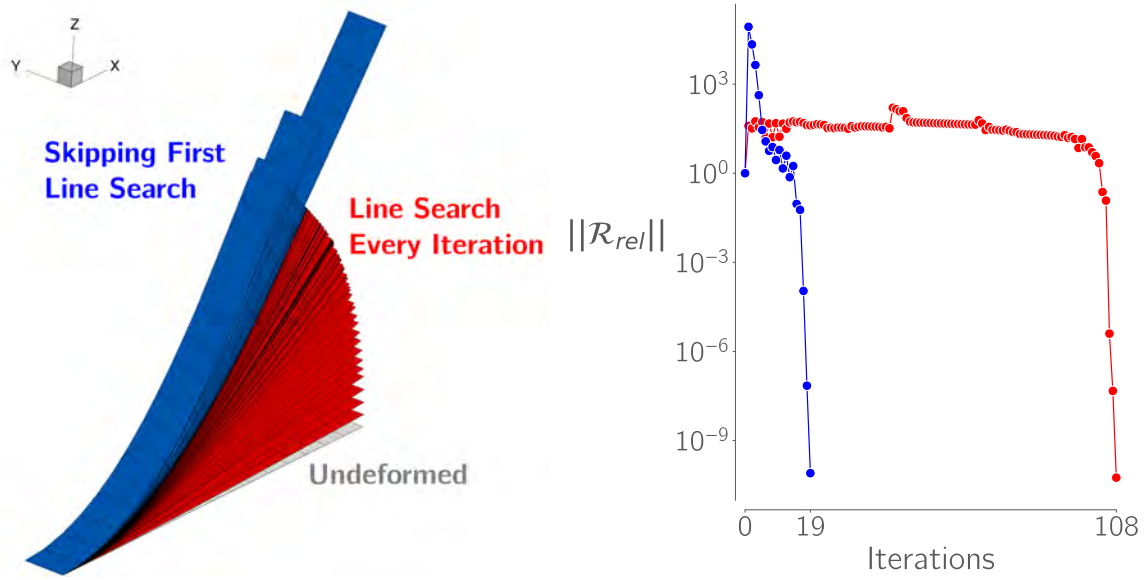

(b) Beam deformation histories.

(c) Residual norm histories.

Figure 3.3: Skipping line searches on the first iteration of each load increment dramatically speeds up nonlinear solution.

I remedy this problem by forgoing any line searches on the first iteration of each load increment and, in subsequent iterations, a line search is only performed if the the computed Newton-Raphson step reduces the relevant line search merit function by less than 50%. As shown in figures 3.3b and 3.3c, these changes drastically improve the rate of convergence of the solver. In the simple cantilever problem demonstrated here, after skipping the line search on the first iteration, the majority of subsequent iterations required no line search, achieving Newton-like convergence and requiring only 19 iterations.

3.1.6. BFGS Updates

Since updating the tangent stiffness matrix is generally the most expensive operation in a nonlinear FE solution process, one of the main ways in which researchers have sought to reduce the computational cost of nonlinear FE solvers is to develop strategies which minimise the number of full order stiffness updates performed. I implement approximate stiffness updates using the BFGS method as proposed by Matthies and Strang [102]. Crucially, rather than explicitly updating the tangent stiffness matrix itself, their method involves only vector-vector products, matrix-vector products and one linear solve with the original stiffness matrix which has already been factored, all of which are considerably cheaper than updating and refactoring the stiffness matrix:

$$\Delta u_i = \left(\left[\mathcal{I}\right] + \left\{w_n\right\}\left\{v_n\right\}^T\right) \dots \left(\left[\mathcal{I}\right] + \left\{w_1\right\}\left\{v_1\right\}^T\right) \left[K_{T,i-n}\right]^{-1} \left(\left[\mathcal{I}\right] + \left\{v_1\right\}\left\{w_1\right\}^T\right) \dots \left(\left[\mathcal{I}\right] + \left\{v_n\right\}\left\{w_n\right\}^T\right) \left\{-\mathcal{R}\left(u_i\right)\right\} \quad (3.16)$$

In testing however, the method, as implemented, proved to have only negative effects on solver performance, struggling to solve even the basic cantilever problem shown above, even with the assistance of line searches. As such, BFGS updates are not used in the remainder of this work.

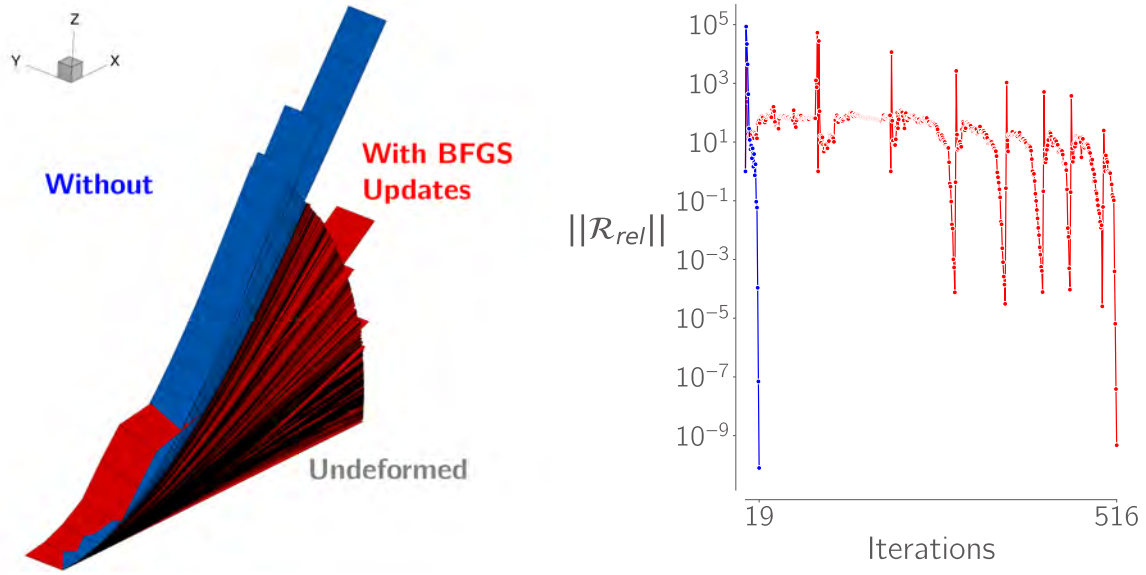

(a) Beam deformation histories.

(b) Residual norm histories.

Figure 3.4: Approximate stiffness updates with the BFGS method result in poor solver stability, resulting in drastically longer solution times.

3.1.7. Solution Restarting

During optimisation or coupled aerostructural analysis, the structural solver is called repeatedly, with slightly different external and/or internal loads. In these cases, the previously converged state, u^*, no longer solves the system of equations but can be used as a good starting point to drastically speed up the next solution. The problem in this case is to decide at what load factor to restart the load incrementation process at. Restarting at $\lambda = 0$ would likely waste the useful initial guess u^* but, if the design or loading has changed significantly, it may not be possible to converge the problem starting at the full load factor, $\lambda = 1$.

To address this problem I present two methods for computing an 'optimal' restart load factor, λ^*, both of which find the load factor that minimises some measure of the solution error. The first method I implement

finds the load factor that minimises the residual norm, a method first proposed by Bergan [103]. Figure 3.5 shows a simple graphical example of the definition of the residual between an internal and scaled external force vector.

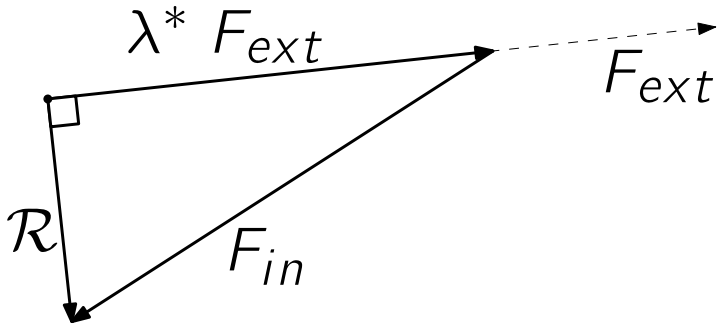

Figure 3.5: A geometric representation of the minimum residual restart method.

Viewing the problem this way makes it clear that the magnitude of the residual is minimised when the load factor is such that the residual is orthogonal to the external force vector. The residual minimising load factor can then be computed as:

$$\left\{F_{ex}\right\}^{\mathrm{T}}\left\{\mathcal{R}\right\} = \left\{F_{ex}\right\}^{\mathrm{T}}\left(\lambda\left\{F_{ex}\right\} + \left\{F_{in}(u^*)\right\}\right) = 0 \Rightarrow \lambda = \frac{-\left\{F_{ex}\right\}^{\mathrm{T}}\left\{F_{in}\right\}}{||F_{ex}||^2} \tag{3.17}$$

In shell problems, terms representing the in-plane, out-of-plane and rotational residuals can have drastically different scales, which can make a Euclidean norm a poor indicator of true magnitude of solution error. To redress this problem, many use an energy based measure of error. In this spirit, I propose a novel method that minimises the strain energy error, defined as the product of the residual and the displacement step resulting from that residual:

$$\Delta E = \left\{R\right\}^{\mathrm{T}}\left[K_T\right]^{-1}\left\{R\right\} = \left\{\left\{F_{in}\right\} + \lambda\left\{F_{ex}\right\}\right\}^{\mathrm{T}}\left[K_T\right]^{-1}\left\{\left\{F_{in}\right\} + \lambda^*\left\{F_{ex}\right\}\right\} \tag{3.18}$$

Differentiating this expression with respect to the load factor and solving for the minimum energy error gives:

$$\lambda^* = \frac{-\left(\left\{\Delta u_i\right\}^{\mathrm{T}}\left\{F_{ex}\right\} + \left\{\Delta u_e\right\}^{\mathrm{T}}\left\{F_{in}\right\}\right)}{2\left\{\Delta u_e\right\}^{\mathrm{T}}\left\{F_{ex}\right\}} \tag{3.19}$$

$$\left\{\Delta u_i\right\} = \left[K_T\right]^{-1}\left\{F_{in}\right\},\ \left\{\Delta u_e\right\} = \left[K_T\right]^{-1}\left\{F_{ex}\right\} \tag{3.20}$$

This minimum energy restart method does require the expensive factorisation of the current tangent stiffness matrix but, as explained previously, TACS direct solution method means that this factorisation can be reused to compute both Δu_i and Δu_e before being reused again for the first iteration of the restarted solution, meaning that the method does not add a significant cost to the solution process.

If the optimal load factor is closer to the final desired load factor than the user defined initial increment size then it is used as the initial load factor for the solution process. If the first load increment in the restarted solution fails to converge the solution is reset and started from the original initial load factor. In the case the the external loads have reduced or the structure has been made stiffer since the previous solution, the optimal load factor may be greater than 1, in which case the direction of load incrementation is reversed.

3.1.8. Solver validation

Before moving onto any analysis studies, I validate my solver and TACS' nonlinear formulation on a series of geometrically nonlinear shell problems defined from the NAFEMS 'Assembly Benchmark Tests for 3-D Beams and Shells Exhibiting Geometric Non-Linear Behaviour' document. Figure 3.6 shows the setup and results from the first benchmark problem, a simple tip loaded cantilever beam, with the TACS solutions compared to analytical results and results from the commercial FE code Abaqus, taken from solutions by Sze et al. [104]. Appendix B contains more details and results of all the benchmarks validated against. Overall, the TACS solutions match both analytical and Abaqus solutions to within 1% for all problems tested.

For successful optimisation, the adjoint gradients computed by TACS must also be accurate. The process of computing adjoint gradients of nonlinear FE formulations is identical to the process for linear formulations, save for ensuring that the tangent stiffness matrix used is up to date with the converged structural state. To validate their accuracy, I compute the adjoint gradients of a variety of TACS output functions with respect to the thickness design variable of the cantilever beam in figure 3.6 with linear and nonlinear element formulations. I then compare these against the same gradients computed via the complex step method which, for the step size used in this study (10^{-100}), can be considered the ground truth. Figure 3.7 shows that the adjoints in the nonlinear case match or exceed the accuracy of the linear formulation adjoints. The only function whose gradients are not accurate is the MaxFailure function which is expected as the maximum stress value in the structure is a discontinuous function.

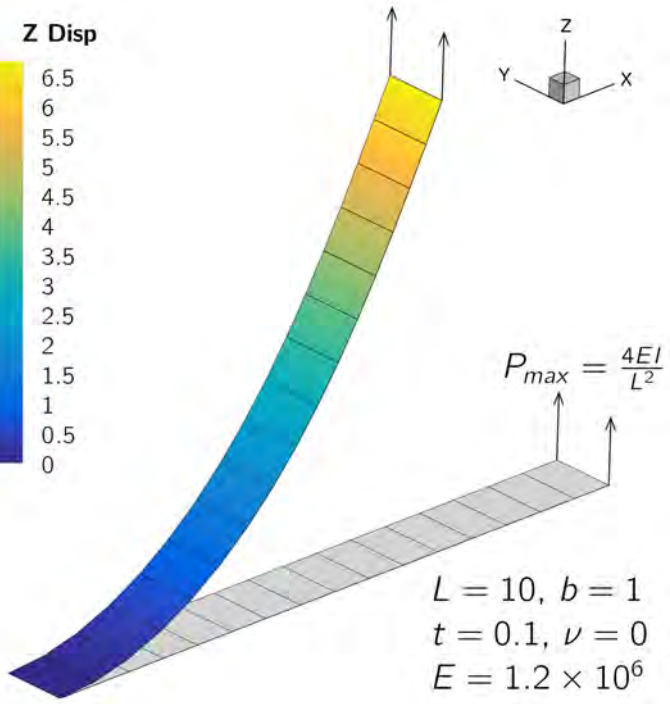

(a) Definition of the benchmark tip loaded cantilever problem from [104]

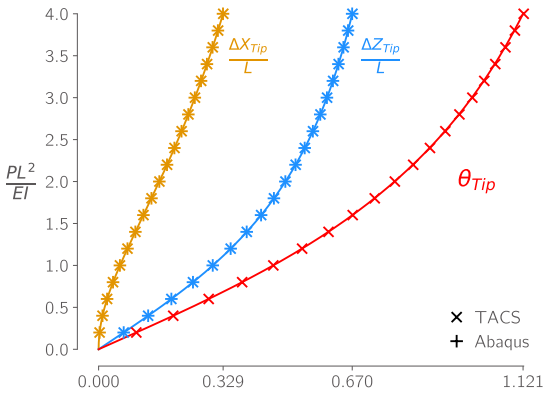

(b) Equilibrium path of the cantilever tip compared with results from theory and Abaqus

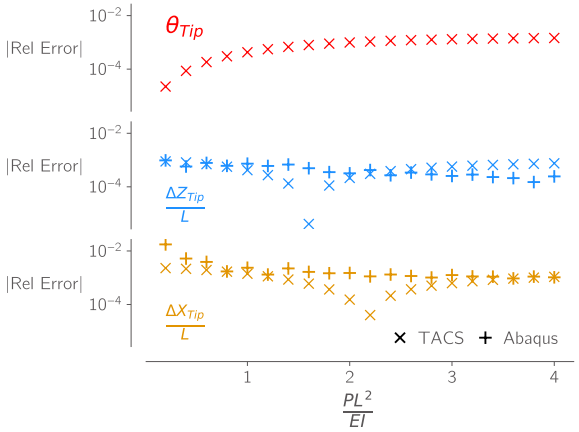

(c) Relative errors of TACS and Abaqus vs theory

Figure 3.6: Validation results for benchmark cantilever beam problem.

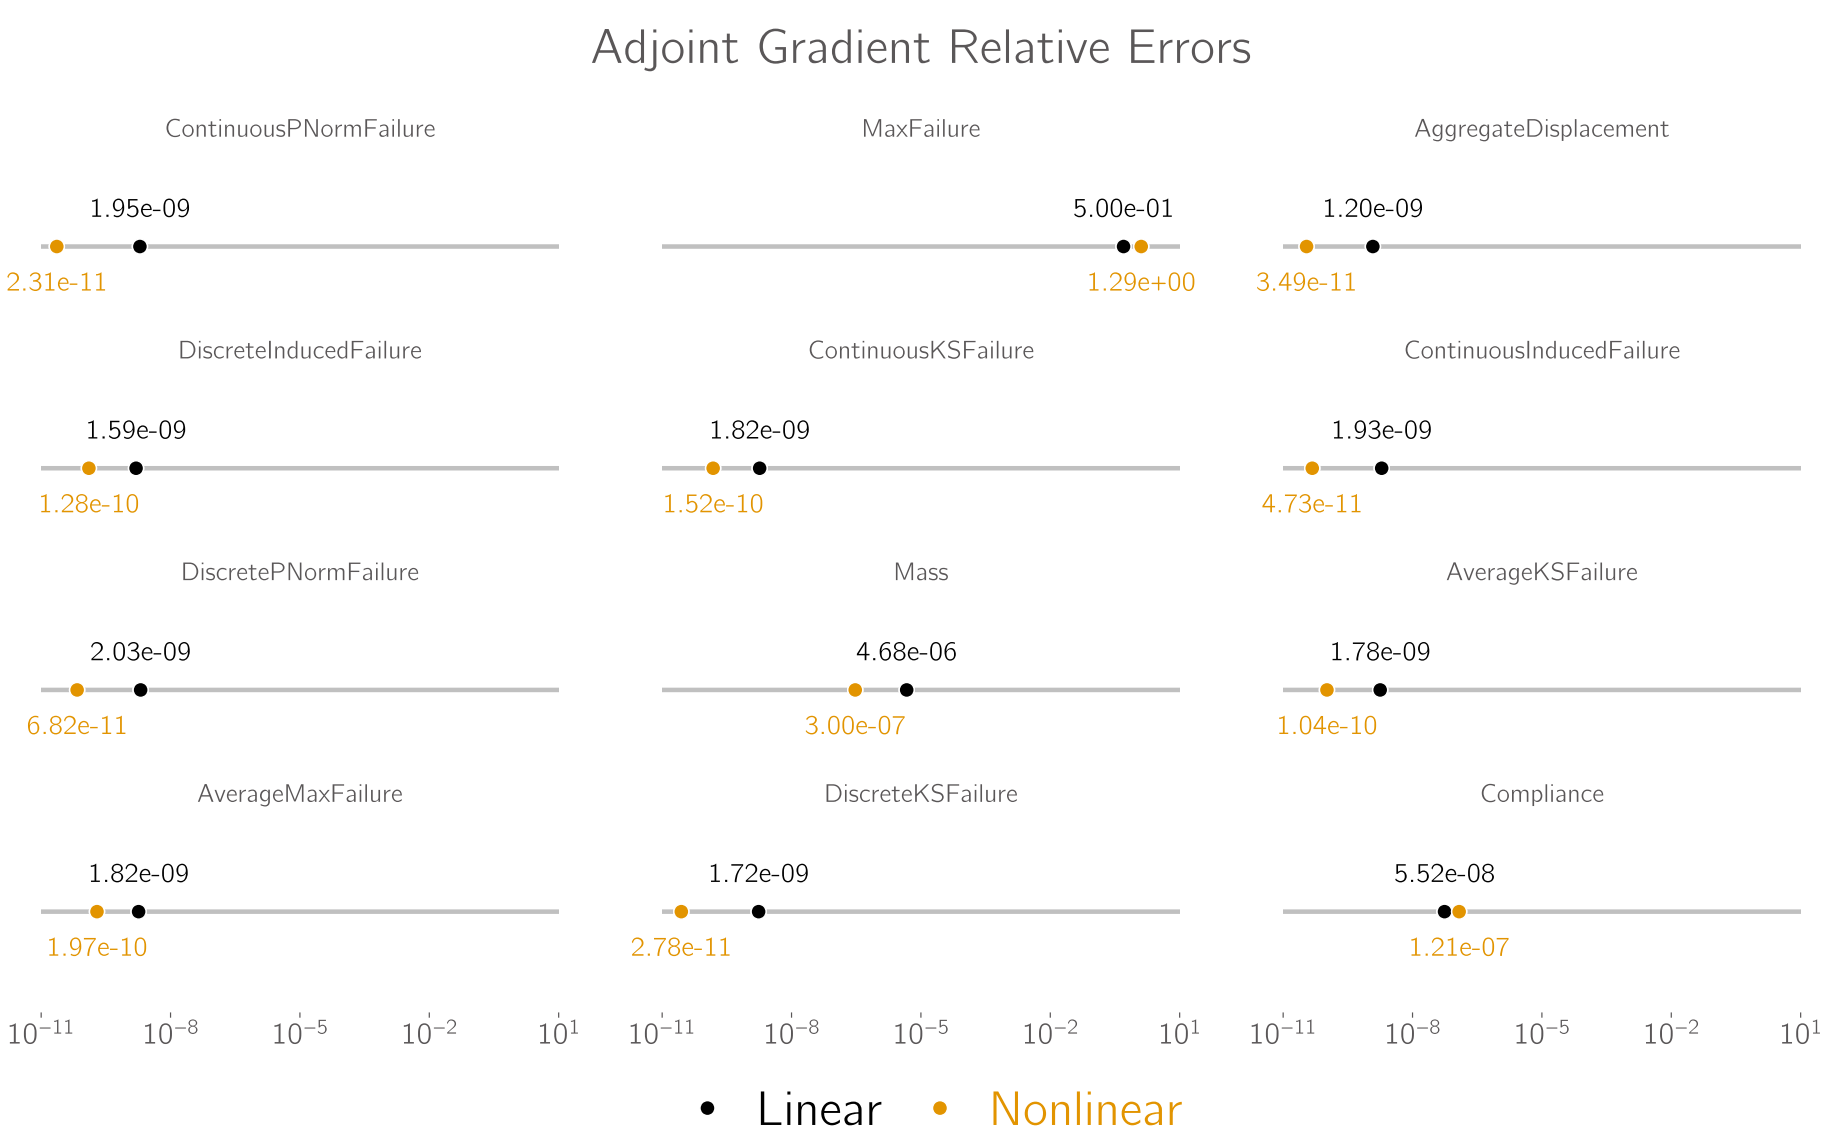

Figure 3.7: When compared to reference values obtained using the complex step method, the adjoint gradients for nonlinear elements match or exceed accuracy of linear elements

3.1.9. Solver Performance Testing

To test which combinations of solver settings result in the best performance, I run a series of analyses on two different wingboxes. For each type of line search I run the same analysis whilst sweeping through different combinations of N_{des} and initial load increment size.

The first model I used for this study is a relatively low aspect ratio wingbox, based approximately on a Boeing 717. The wingbox is subject to a large uniform pressure difference of 140 kPa between the upper and lower skins, resulting in a tip displacement of roughly 50% of the wing's 14 m semi-span, making it a very challenging case to converge. The second case is an analysis of the uCRM-13.5 under 2.5 g manoeuvre, the details of which are described in more detail in the next chapter. The deformed and undeformed shapes of each wing are shown in figure 3.8. The Boeing 717 analyses are run on 8 cores, resulting in around 1000 elements per core whilst the uCRM analyses are run on 16 cores, with approximately 1500 elements per core.

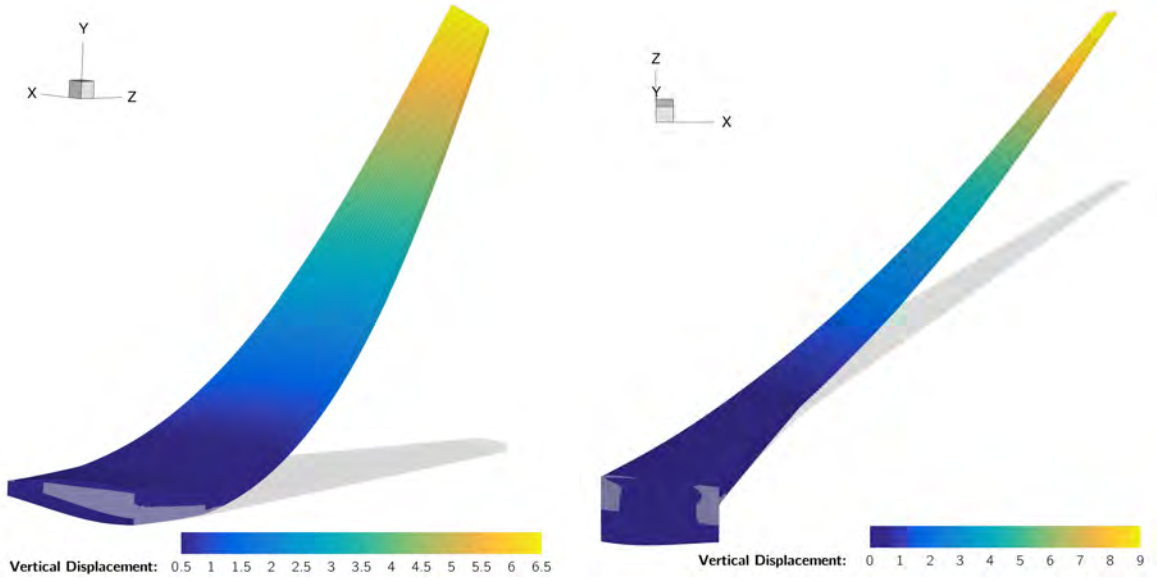

(a) The Boeing 717 wingbox model.

(b) The uCRM-13.5 wingbox model.

Figure 3.8: The two cases used for solver performance testing.

Figures 3.9 and 3.10 show the solution times in seconds for all combinations of line search methods and load incrementation settings tested. In the challenging Boeing 717 wing case, for all line search types, the smallest tested value of desired iterations, 4, results in the lowest solution times, indicating the most efficient strategy for this case requires maintaining very small load steps throughout the solution process. Additionally, for all the residual based line searches, the smallest initial load increment, in conjunction with a desired iteration value of 4 resulted in the shortest solution time, this suggests that the residual based methods are unable converge the larger initial increments, repeatedly diverging or reaching the increment iteration limit before reaching a small enough load increment, presumably around 0.01. In contrast, the minimum energy line search method achieves it's minimum solution time with an initial increment size of 0.1, proving the method's superior convergence rate compared to the residual based methods. Indeed, over the entire spectrum of incrementation settings, the energy based similarly performs significantly better than the residual based methods.

The uCRM case proves a much more stable and simpler problem to solve, as indicated in part by the solution times which are generally 1-2 orders of magnitude lower than for the Boeing case. This stability is also indicated by the fact that, for all line search types, choosing the smallest initial increment and the lowest desired iterations results in the longest solution times while increasing the initial increment size and desired number of iterations reliably results in lower solve times. Even the energy based line search method, however, is still slower than the method with no line search for most combinations of settings.

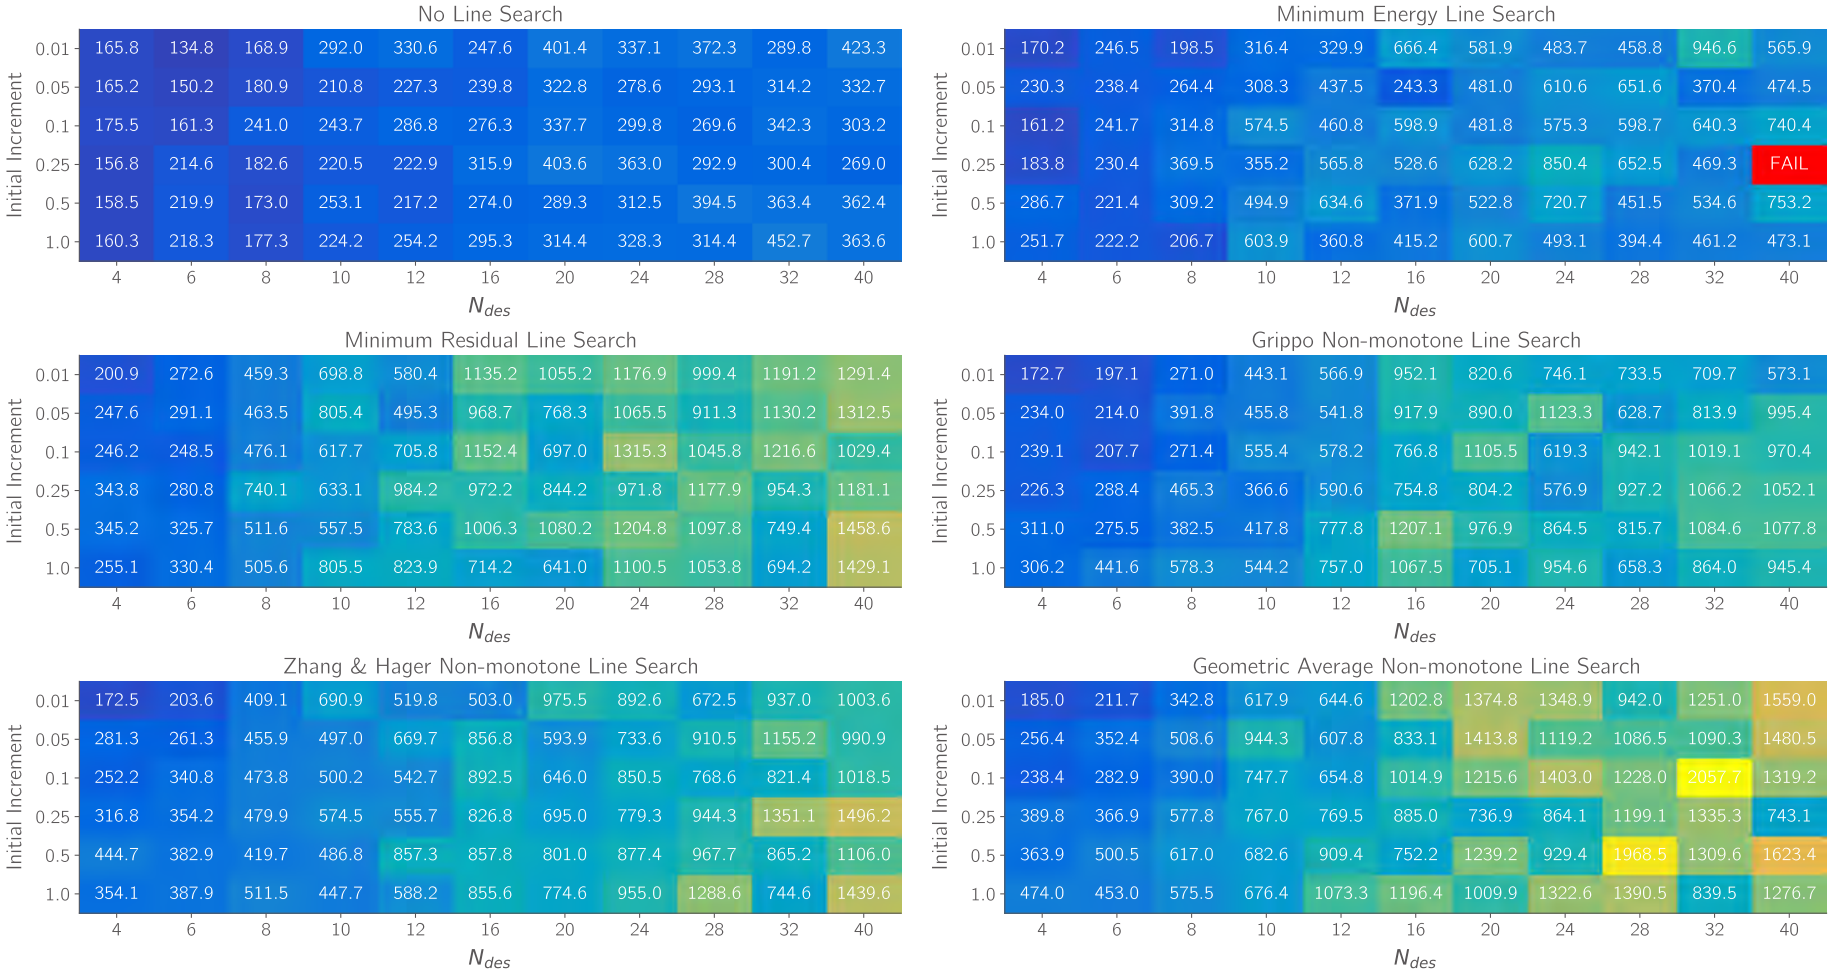

Figure 3.9: Analysis times for the B717 wingbox

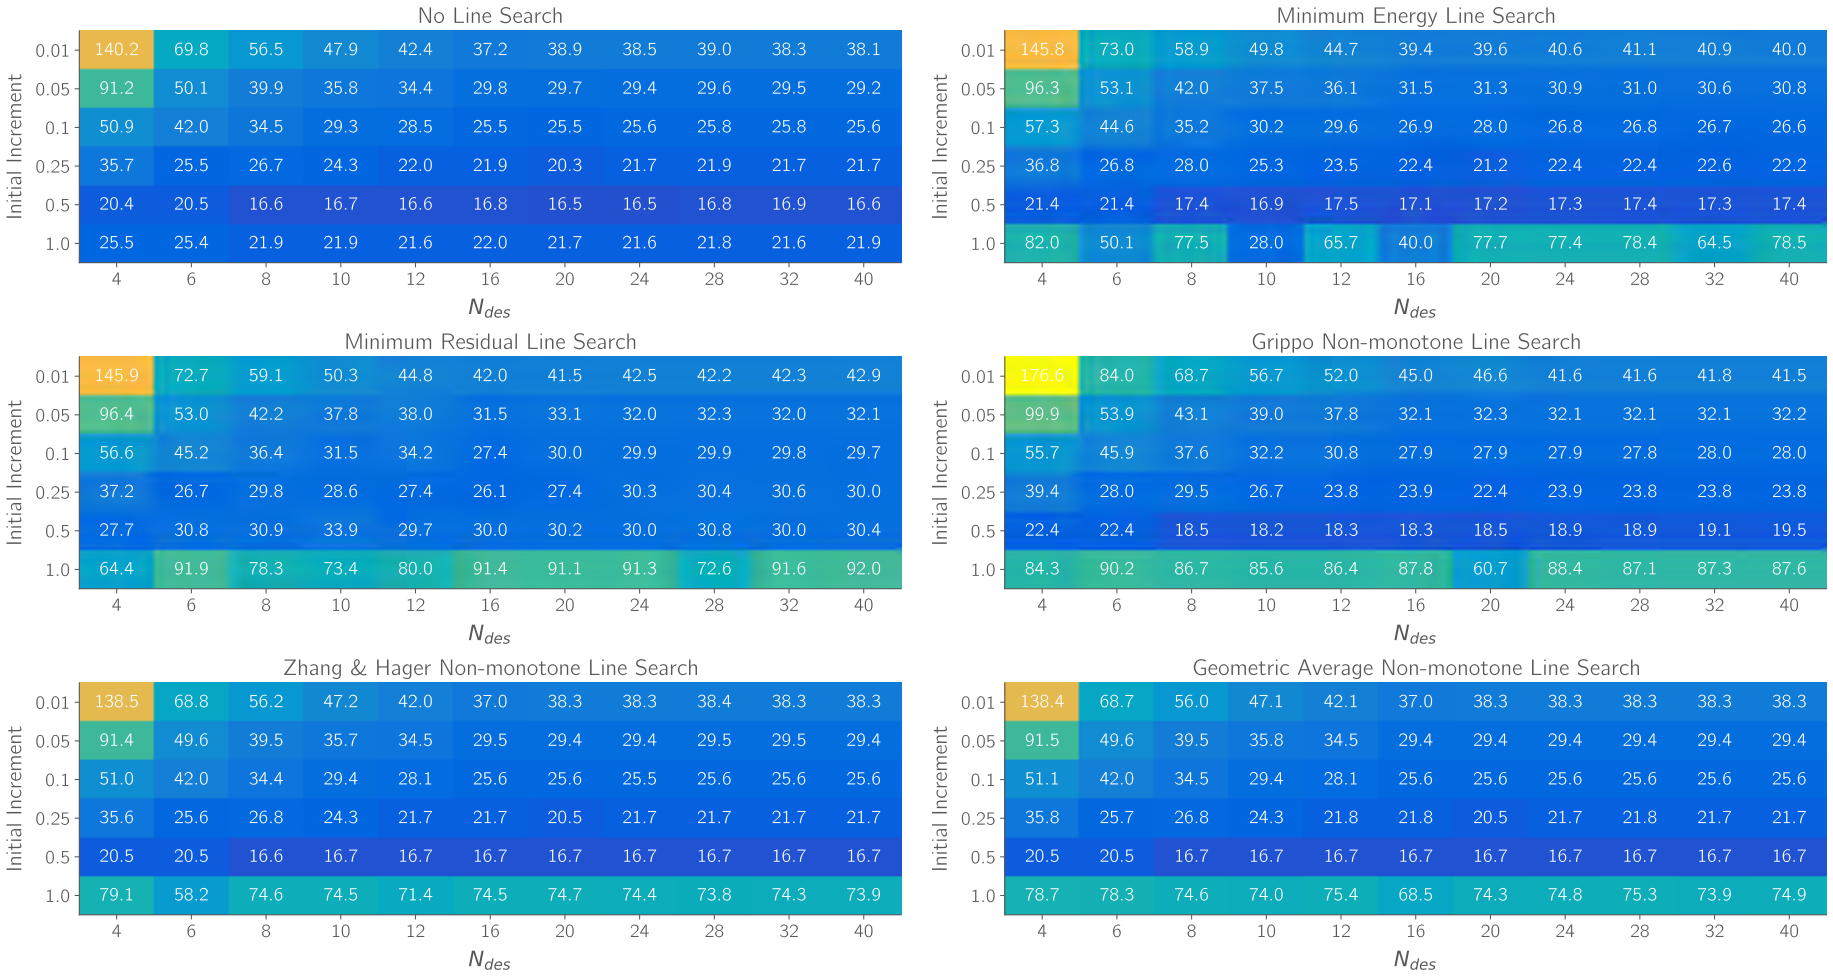

Figure 3.10: Analysis times for the uCRM-13.5 wingbox

Solver Performance Conclusions

From these results I draw 2 primary conclusions:

It's hard to beat a straight-forward Newton-Raphson method: All of the implemented line search methods have the intended effect of improving the stability of the Newton-Raphson solver, but at the cost of much slower solutions. The energy based line search method of Matthies and Strang [102] stands out as the best of the implemented methods, able to converge larger loadsteps than the residual based methods on practical timescales. However, although for a handful of combinations of settings the energy based is fastest, the base Newton-Raphson method with no line search is fastest for the vast majority of cases tested and achieves the lowest solution time for both test cases.

Load incrementation is the key driver of solver performance: The greatest influence seen on the solution time of the Newton-Raphson solver comes from the load incrementation strategy, the aim of which is to maximise the rate at which the load factor is increased while maintaining fast and consistent convergence in each load increment. For challenging problems, like the Boeing 717 test case, a very conservative load incrementation strategy, with a small initial load increment and low N_{des}, is key to keeping solver within the basin of attraction of the solution and achieving good solution times. Performance is most sensitive to N_{des}. For more stable problems like the uCRM test case, where only a few load increments are required, the solver has fewer opportunities to adpat the load incrementation rate and solution time is therefore more sensitive to the size of the initial load increment than to N_{des}.

3.2. Aerostructural Solver Background

Aerostructural or aeroelastic analysis is a subset of the wider field of fluid-structure interaction (FSI), which studies the coupled behaviour of fluid flows and movable bodies. From a computational point of view, solving an FSI problem involves satisfying at least two sets of residual equations:

$$\left\{\mathcal{R}_f(u_f, X_f(u_s))\right\} = 0 \tag{3.21}$$

$$\left\{\mathcal{R}_s(u_s, F_s(u_f))\right\} = 0 \tag{3.22}$$

Where the $\mathcal{R}_f$ and u_f are the fluid residuals and states, and $\mathcal{R}_s$ and u_s are the structural residuals and states. The coupling of the two systems of equations occurs through the displacements in the boundaries of the fluid problem, X_f, which are dependent on the motion of the structure, and the external forces acting on the structure, F_s, which are dependent on the fluid states. These coupling of these variables are described by force and displacement transfer schemes, T_F and T_X:

$$\left\{F_s\right\} = T_F(u_f) \tag{3.23}$$

$$\left\{X_f\right\} = T_X(u_s) \tag{3.24}$$

Methods for solving these coupled sets of equations can be classified as either partitioned or monolithic, each with their own benefits and drawbacks. A partitioned approach involves solving the fluid and structural problems using separate solvers, occasionally exchanging information between the two in order to converge to a consistent solution. In a monolithic approach, the fluid and structural problems are combined into a single system of equations and solved with a single solver. This makes monolithic solvers potentially more robust and efficient than partitioned approaches, particularly for strongly coupled problems (e.g analysing flexible aircraft at high load factors). However these benefits is significantly more difficult for a number of reasons:

1. A new solver must be developed for the coupled system rather than re-using the individual solvers for each discipline

2. Monolithic solvers requires more in depth access to the states of each system than may be available from commercial codes.
3. Specialised preconditioning is required for good performance in the solution of the coupled linear system because it is typically very large, badly conditioned, and unstructured [105].
4. When using CFD, each monolithic iteration requires the CFD volume mesh to be deformed and so the approach requires an efficient mesh deformation algorithm.
5. The off-diagonal blocks in the coupled Jacobian matrix require the derivatives of the load and displacement transfer schemes, which requires additional implementation.

MACH contains both partitioned and monolithic solvers for aerostructural problems [6].

Algorithm 1 presents a simplified version of MACH's partitioned aerostructural solver which is used throughout this thesis. The solver employs a Gauss-Seidel iteration scheme, meaning that the CFD and FE solvers are run sequentially with forces or displacements being transferred between solvers after each solve. An alternative approach is the Jacobi iteration scheme, in which both solvers are run simultaneously, with forces and displacements being transferred after the completion of both solvers. Despite the potential speed-up due to greater parallelisation, this scheme is generally suffers from poorer stability and slower convergence as each solver is working with out of date information about the state of the other discipline.

An under-relaxation factor θ is used to control the magnitude the displacement update applied in each iteration. This under-relaxation factor improves the coupled solver's performance by reducing over-shooting and can be fixed or chosen adaptively using Aitken acceleration [106]. Kenway et al. [6] found that the best choice of fixed under-relaxation factor lead to slightly improved performance compared to Aitken acceleration but cannot generally be predicted a-priori. At increased load factors, the optimum fixed under-relaxation factor reduces and the performance benefit over the adaptively chosen factor reduce.

Algorithm 1: MACH's Gauss-Seidel partitioned aerostructural solver [6]

Given: $\left\{u_s^{(0)}\right\}, \left\{u_f^{(0)}\right\}, \theta^{(1)}$;
while *Not converged* **do**
  $\left\{X_f^{(n)}\right\} = T_X(u_s^{(n-1)})$; ▷ Transfer structure displacements and deform aero mesh
  Find $\left\{u_f^{(n)}\right\}$ $s.t$ $\left\|\mathcal{R}_f\left(u_f^{(n)}, X_f^{(n)}\right)\right\| \leq \delta_{f,rel}\left\|\mathcal{R}_f\left(u_f^{(n-1)}, X_f^{(n)}\right)\right\|$; ▷ Approximately solve aero
  $\left\{F_s^{(n)}\right\} = T_F(u_f^{(n)})$; ▷ Transfer aero forces to structure
  Find $\left\{u_s^{(n)}\right\}$ $s.t$ $\left\|\mathcal{R}_s\left(u_s^{(n)}, F_s^{(n)}\right)\right\| \leq \delta_{s,rel}\|F_{ex}\|$; ▷ Solve structure
  $\left\{\Delta u_s^n\right\} = \left\{u_s^{(n)}\right\} - \left\{u_s^{(n-1)}\right\}$;
  if $n > 1$ **then**
    $\theta^{(n)} = 1 - \frac{\left\{\Delta u_s^{(n)} - \Delta u_s^{(n-1)}\right\}^T \left\{\Delta u_s^{(n)}\right\}}{\left\|\Delta u_s^{(n)} - \Delta u_s^{(n-1)}\right\|^2}$; ▷ Update under-relaxation factor using Aitken acceleration
  end
  $\left\{u_s^{(n)}\right\} = \left\{u_s^{(n)}\right\} + \theta^{(n)}\left\{\Delta u_s^{(n)}\right\}$; ▷ Apply under-relaxed displacement increment
end

3.3. Aerostructural Solver Implementation

The FE solver described in the previous section is implemented such that it can be accessed through the same calls to the pyTACS interface as the linear solver. As such, no modifications to the source of pyAeroStruct were strictly necessary in order to perform geometrically nonlinear aerostructural analyses.

I do however implement two new features which aim to reduce solution times when using a nonlinear structural formulation.

First, because nonlinear systems do not abide by the principle of linear superposition, the structural state resulting from the aerostructural solver's displacement under-relaxation, $\left\{u_s\right\}$, is typically not close to a solution for any scaling of the applied loading. It is therefore typically not efficient, or even possible, to restart the subsequent structural analysis from this damped state until late in the aerostructural solution process when displacement updates become very small. To remedy this, I save the converged solution of each structural analysis before damping, $\left\{\hat{u}_s\right\}$, and restart the next structural analysis from this undamped state.

The second feature stems from the fact that in the initial iterations of the aerostructural solution process, the aerodynamic forces are often drastically overpredicted, potentially by a factor of 2-3. This is due in part to the partial convergence of the CFD solver and the lack of any load alleviation of the wing in its undeformed state. For a linear structural analysis, this poses no problems as the cost of solving the pre-factorised linear system is negligible and independent on the magnitude of the loading. However, assuming a consistent load incrementation rate, the time to solve a nonlinear structural analysis is approximately linearly proportional to the loading magnitude, meaning the first few structural analyses may take 2-3 times longer than usual. In the worst case, it may not be possible to solve the problem at all if unrealistically high aerodynamic forces result in explicit buckling of the FE model.

To avoid this potential issue and hopefully improve the speed of the initial structural analyses, I implement the load ramping method proposed by Sanchez et al. [107]. This is achieved by ramping the maximum load scale to which the structural solver increments over the first N Gauss-Seidel iterations following a smooth polynomial ramping function:

$$\lambda_{ramp} = -2\left(\frac{n}{N}\right)^3 + 3\left(\frac{n}{N}\right)^2 \tag{3.25}$$

A fixed under-relaxation factor is used during the load ramp before resuming Aitken acceleration once the ramp is complete. To enable efficient restarting of the aerostructural solver without restarting the load ramp from zero, I compute the optimal restart load factor using the same minimum strain energy method described earlier before each call of the structural solver. This value is then clipped to lie between the current load ramp value and 1 before being used as the maximum load scale for the structural analysis. The modified solution algorithm is described in algorithm 2.

To test the effectiveness of this load ramping strategy, I run aerostructural analyses of the Boeing 717 wing used in the structural solver performance tests with and without load ramping. Some of the key values describing the test case are shown in table 3.1. To make the problem more challenging, I reduce the elastic modulus of the wingbox to 3 GPa to achieve a large vertical displacement of 22% semispan and to increase the nonlinearity in the aerodynamics I use a transonic Mach number and a relatively high angle of attack which results in significant flow separation over the wing in the early stages of the solution.

Table 3.1: Key values from the aerostructural solver performance test.

Parameter	**Value**	**Unit**
Altitude	10000	m
α	8	°
Mach	0.85	
Tip displacement	≈22	% Semispan

Figure 3.12 shows the time for aerostructural solutions with no load ramping ($N = 1$) and load ramps of varying length from 2 to 10 iterations. The load ramping results in a lower solution time in all cases but

Algorithm 2: MACH's Gauss-Seidel partitioned aerostructural solver with additions for geometrically nonlinear analysis highlighted.

Given: $\left\{u_s^{(0)}\right\}, \left\{u_f^{(0)}\right\}, \theta^{(1)}$

while *Not converged* **do**

- $\left\{X_f^{(n)}\right\} = T_X(u_s^{(n-1)})$ ▷ Transfer structure displacements and deform aero mesh
- Find $\left\{u_f^{(n)}\right\}$ $s.t.$ $\left\|\mathcal{R}_f\left(u_f^{(n)}, X_f^{(n)}\right)\right\| \leq \delta_{f,rel}\left\|\mathcal{R}_f\left(u_f^{(n-1)}, X_f^{(n)}\right)\right\|$ ▷ Approximately solve aero
- $\left\{F_s^{(n)}\right\} = T_F(u_f^{(n)})$ ▷ Transfer aerodynamic forces to structure
- **if** *Ramping* **then**
  - $\lambda^* = \dfrac{-\left(\left\{\Delta u_i\right\}^T\left\{F_{ex}\right\}+\left\{\Delta u_e\right\}^T\left\{F_{in}\right\}\right)}{2\left\{\Delta u_e\right\}^T\left\{F_{ex}\right\}}$ ▷ Compute optimal restart load factor
  - **if** $\lambda^* > 0.95$ **then**
    - $\lambda^* = 1$
    - Ramping = False ▷ End ramping if λ^* close to unity
  - **end**
  - $\lambda_{ramp} = -2\left(\frac{n}{N}\right)^3 + 3\left(\frac{n}{N}\right)^2$ ▷ Compute ramp load factor
  - $\lambda^{(n)} = \max\left(\lambda_{ramp}, \lambda^*\right)$ ▷ Choose maximum load factor
- **else**
  - $\lambda^{(n)} = 1$
- **end**
- Find $\left\{\hat{u}_s^{(n)}\right\}$ $s.t.$ $\left\|\mathcal{R}_s\left(u_s^{(n)}, \lambda^{(n)} F_s^{(n)}\right)\right\| \leq \delta_{s,rel}\lambda^{(n)}\left\|F_{ex}\right\|$ ▷ Solve structure with scaled load
- $\left\{\Delta u_s^{(n)}\right\} = \left\{\hat{u}_s^{(n)}\right\} - \left\{u_s^{(n-1)}\right\}$
- **if** $n > 1$ **then**
  - **if** *Ramping* **then**
    - $\theta^{(n)} = \theta^{(n-1)}$
  - **else**
    - $\theta^{(n)} = 1 - \dfrac{\left\{\Delta u_s^{(n)} - \Delta u_s^{(n-1)}\right\}^T\left\{\Delta u_s^{(n)}\right\}}{\left\|\Delta u_s^{(n)} - \Delta u_s^{(n-1)}\right\|^2}$ ▷ Update under-relaxation factor using Aitken acceleration
  - **end**
- **end**
- $\left\{u_s^{(n)}\right\} = \left\{u_s^{(n)}\right\} + \theta^{(n)}\left\{\Delta u_s^{(n)}\right\}$ ▷ Apply under-relaxed displacement increment

end

there is no clear trend with both short ($N = 2$) and long ($N = 8$) ramps achieving a 40% speed-up over a non-ramped solution.

Figure 3.13 shows the convergence histories from a selection of the analyses, showing the norms of the aerodynamic and structural residuals and the structural displacement vector. The effect of the load ramping is clear in the early stages of the displacement norm. The non-ramped case is slower to converge as the initial structural analyses are solved to the full load scale, which takes significantly longer. The first iterations of the ramped cases take far less time than the non-ramped case and the wing displacement magnitude increases gradually over the ramp, whereas the displacement in the non-ramped case shows significant oscillations. These large oscillations in the deformed wing shape further slow down the solution as they result in large changes in the flow around the wing. This not only causes the aerodynamic residual to converge more slowly, but can also result in the structural solution restarting failing, requiring a total

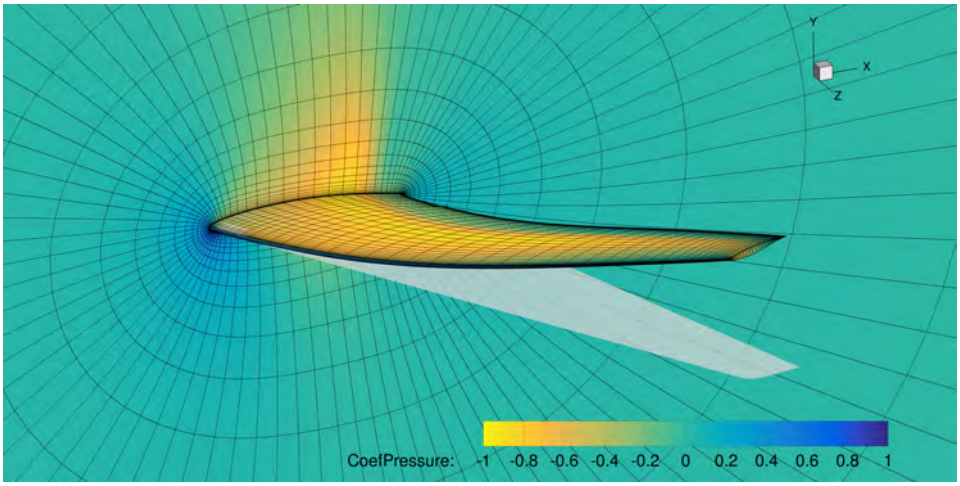

Figure 3.11: The Boeing 717 wing case used for the aerostructural load ramping test.

reset of the structural solution.

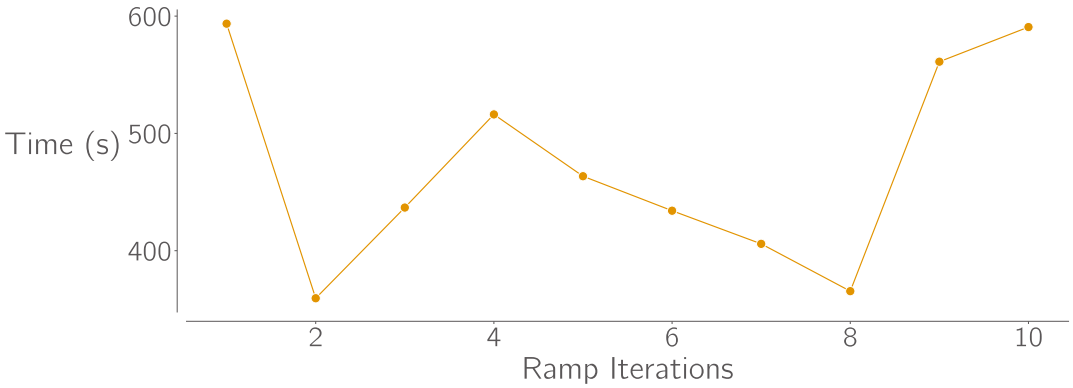

Figure 3.12: Solution time vs load ramp length for the Boeing 717 test case. All load ramps result in a solution speed-up but with no clear trend.

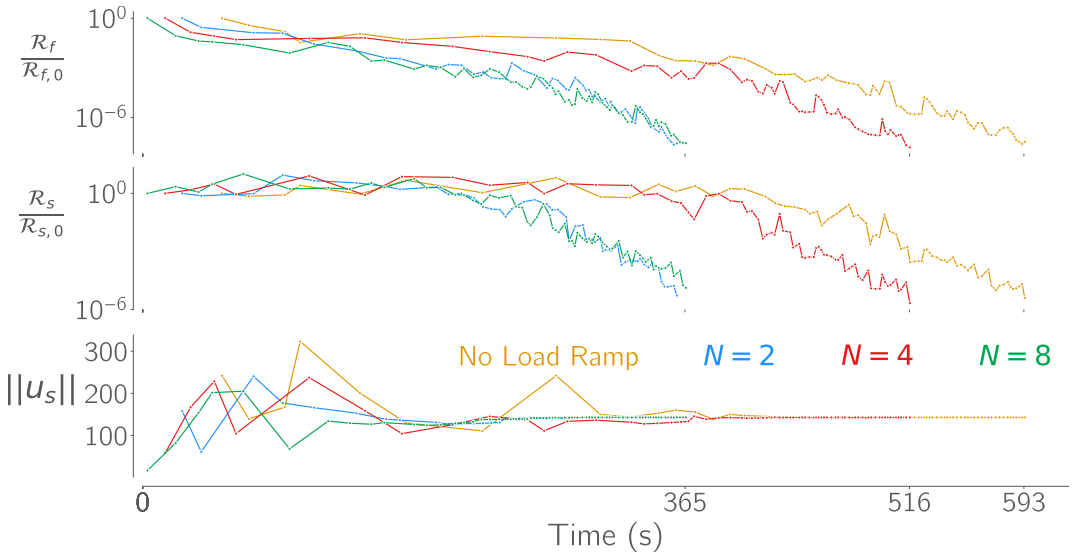

Figure 3.13: Both short ($N = 2$) and long ($N = 8$) load ramping stages can achieve a 40% speedup over the non-ramped solution. Most of the gains are made in the early Gauss-Seidel iterations.

3.4. Chapter Summary

In this chapter I presented the implementation and performance testing of solvers for geometrically nonlinear structural and aerostructural analysis.

I implemented a Newton-Raphson based nonlinear solver in the python interface to MACH's high performance finite element library, TACS. In an attempt improve the performance of the solver I implemented multiple line search methods and an adaptive load incrementation strategy. After validating the analysis and gradient computation accuracy of the solver, I performed a large set of performance tests on two different wingbox analysis test cases of differing difficulty. The results of these tests showed that the energy based method of Matthies and Strang [102] greatly outperformed all other line search techniques but was still slower than the baseline Newton-Raphson method. The results also showed that the parameters controlling the adaptive load incrementation, the initial load increment size and the desired number of iterations per increment, strongly affect the performance of the solver. For the more challenging wingbox problem, the fastest solution times were achieved with low values for these two parameters, biasing the solver towards taking many small load steps while in the more stable test case, the opposite was true. I also introduced two methods for computing the optimal load factor from which to restart a nonlinear structural analysis from a previous solution. One based on minimising the structural residual and another novel method based on strain energy minimisation, designed to eliminate issues with the poor scaling of shell structure problems.

I then implemented some minor modifications in MACH's aerostructural Gauss-Seidel solver to improve its peformance when using the geometrically nonlinear structural formulation. The most substantial of these modifications was the addition of a load scaling ramp as originally proposed by Sanchez et al. [107] to reduce the time spent solving unrealistically highly loaded structural analyses in the early Gauss-Seidel iterations. I tested the load ramping strategy on a highly flexible transonic wing test case and showed that this ramping method can reduce time to reach a coupled solution by up to 40%.

4

Analysis Studies

This chapter focuses on comparing the results of structural and aerostructural analysis using TACS' linear and nonlinear formulations. I first describe the undeflected common research model (uCRM)s used for these analyses, providing both the motivation for the use of these particular aircraft and the details of the computational models used to represent them. In section 4.2, I describe and present results from a set of linear and nonlinear structural analyses under fixed loading and discuss the resulting differences in the stresses and deformations in the structures due to purely structural nonlinearity. Next, in section 4.3 I describe a similar study, this time using coupled aerostructural analysis, performed both at constant angles of attack and constant lift. Finally, I consolidate the results of these studies, drawing some conclusions on the nonlinear effects responsible for the differences seen between linear and nonlinear analyses and, looking forward, speculate on the influence of these effects on the results of structural and aerostructural optimisation.

4.1. The uCRM Models

Since 2008, the NASA common research model (CRM) has served as a benchmark aircraft model for the aerospace community to test their CFD solvers and aerodynamic shape optimisation frameworks on in realistic transonic aircraft applications [108, 109]. The uCRMs aim to serve the same the same purpose for coupled aerostructural analysis and optimisation. The models were developed by Brooks et al. [81] and further description of the design process of the models is given in their work. In this section, I present the relevant details of the uCRM geometries and computational models used in the remainder of this thesis.

4.1.1. Geometries

Table 4.1 contains some of the key parameters defining the uCRM wing planforms. The uCRM-9 is designed to be an exact aerostructural replica of the original NASA CRM and thus shares the same planform definition. Brooks et al. [81] created the model using a combined inverse design and structural optimisation procedure, to find the jig shape and structural sizing that result in the same deformed wing shape as the original CRM in a cruise condition whilst also meeting failure criteria in manoeuvre conditions.

The uCRM-13.5 is a high aspect-ratio version of the same aircraft. The wing retains the same reference area, sweep and 1/4 chord mean aerodynamic chord (MAC) position in order to remain feasible with the same fuselage, engine and tail configuration. As shown in figure 4.2, while the uCRM-9 is typical of the aspect-ratios seen in current transport aircraft, the uCRM-13.5 represents the kind of wings which may be present on the next generation of transport aircraft entering service in the next two decades. For this reason, I perform all analysis and optimisation studies in this thesis on both the uCRM-9 and 13.5 in order to study how much the importance of considering geometric nonlinearity in design will increase moving from current to next-generation transport aircraft.

Table 4.1: Key parameters of the uCRM planforms [81].

Parameter	uCRM-9	uCRM-13.5	Unit
Aspect-ratio	9	13.5	
Taper-ratio	0.275	0.25	
Span	58.76	72.0	m
Root chord	11.92	11.07	m
Tip chord	2.736	2.06	m
Reference area	383.74	383.78	m^2
1/4 chord sweep	35	35	°

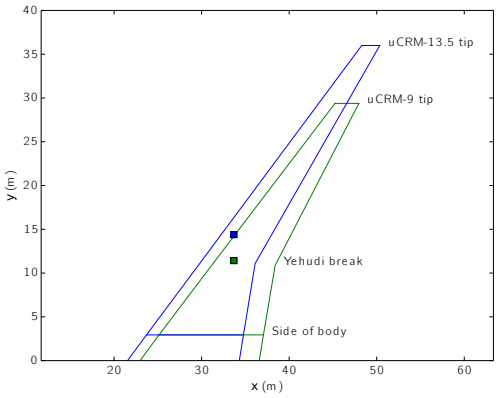

Figure 4.1: Schematics of the uCRM planforms [81]

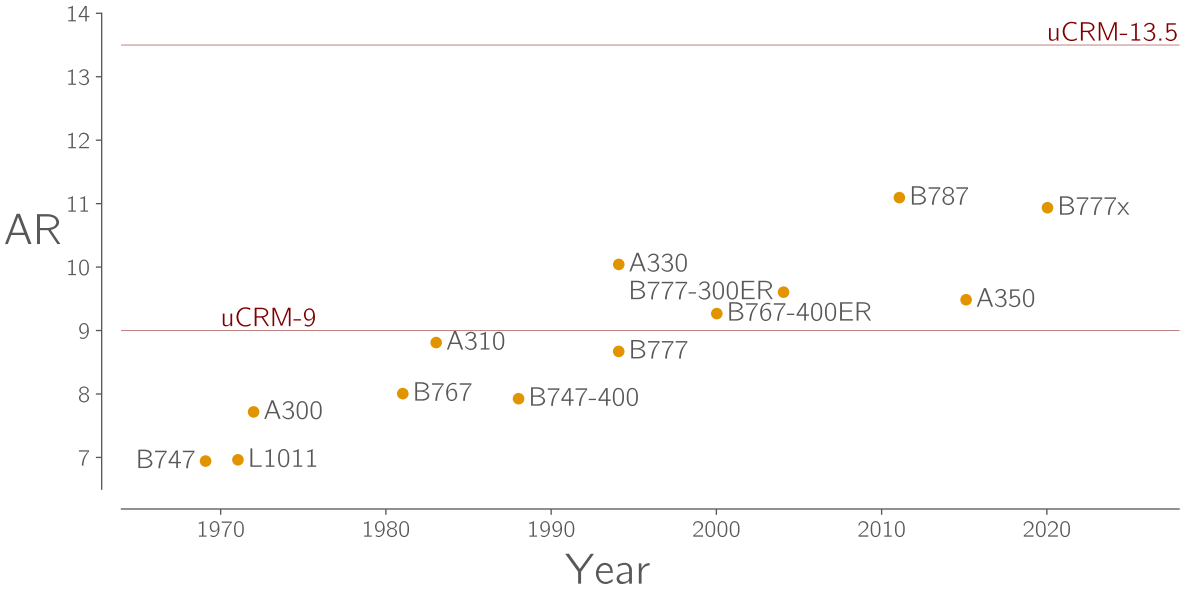

Figure 4.2: The uCRMs have aspect-ratios representing aircraft typically in service today and those which may enter service in the coming decades.

4.1.2. Flight Points

Brooks et al. [81] considered a total of 10 different flight conditions including cruise conditions for computing fuel burn, manoeuvre conditions for constraining structural failure and buffet conditions for constraining buffet onset at high lift and Mach conditions. In this work I consider only the three manoeuvre conditions and, in some cases, a single cruise condition, all of which are listed in table 4.2.

The three manoeuvre cases utilised cover the most extreme positive and negative load factor static manoeuvres that must be sustained by an aircraft of the uCRM's mass and a third 'gust' condition which serves as a proxy for two different structural constraints. Kenway et al. [110] found that stresses in the wingbox became much greater in cruise conditions due to highly effective aeroelastic tailoring when performing aerostructural optimisation considering only the 2.5 and -1 g manoeuvres just mentioned. These increased cruise stresses greatly reduce the aircraft's margin to failure in gust encounters and are detrimental to the fatigue life of the wingbox. In future works, Kenway and Martins [5] and Brooks et al. [81] therefore included a third 1 g manoeuvre condition at the Mach crossover point, where an increased safety factor of 2.67 is applied to the structural failure criteria, compared to the factor of 1.5 used for the other two manoeuvre cases. This increased load factor is based the FAR 23.341 gust load factor formula as described

by Mader et al. [111].

Table 4.2: Flight conditions used in this work.

Condition	Mach	C_L/Lift	Altitude (ft)
Cruise	0.85	0.5	37000
2.5 g manoeuvre	0.64	2.5×MTOW[1]	0
-1 g manoeuvre	0.64	−MTOW[1]	0
1 g 'gust' manoeuvre	0.85	MTOW[1]	27300

[1] maximum take-off weight (MTOW)

4.1.3. Structural Model

Wingbox layout

The wingboxes of the uCRMs share an identical topology, with upper and lower skins, trailing and leading edge spars, ribs oriented perpendicular to the leading edge spar and an additional engine mount panel between the two ribs at the Yehudi break. To maintain the same inter-rib spacing the uCRM-13.5 has 58 ribs while the uCRM-9 has 49. Each rib and each section of the contiguous skins and spars between the ribs is considered a separate panel whose sizing variables can take on their own values. In the skin panels, the longitudinal axis of the smeared blade stiffeners is aligned with the leading edge spar, thus the stiffeners in the skin sections inside the fuselage are not swept whilst those outside the fuselage are. Note that is a departure from the work of Brooks et al. [81], who maintained the same swept stiffener orientation over the entirety of the skins. The stiffeners on the rib and spars are oriented vertically.

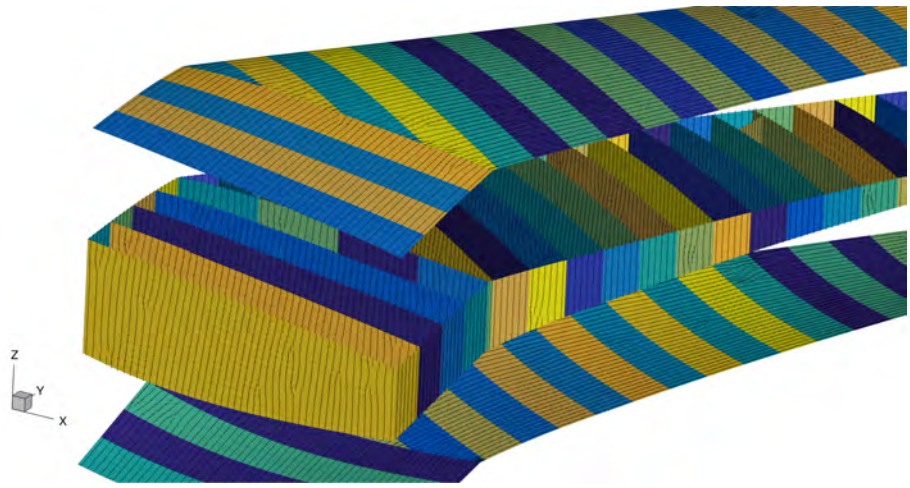

Figure 4.3: The root of the uCRM-9 wingbox with lines indicating stiffener orientation and colours distinguishing separate panels.

Boundary Conditions

Figure 4.4 shows how the uCRM wingboxes are constrained. The rib at the aircraft centreline is clamped while the the rib at the fuselage junction is fixed in the chordwise and vertical directions.

Loading

A variety of different loads are applied to the structural models:

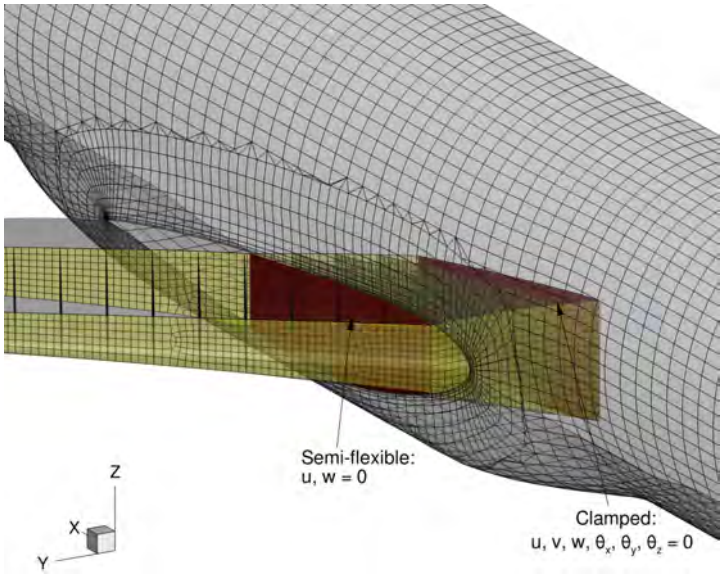

Figure 4.4: The boundary conditions applied to the uCRM wingboxes.

Aerodynamic Loads are applied as nodal forces on the structure using the rigid link load transfer scheme described in chapter 2. During aerostructural analysis, aerodynamic tractions come directly from ADflow, whilst for purely structural analysis they can be loaded into TACS from a file output by a previously run analysis in ADflow or pyAeroStruct.

Gravitational Loads are applied to model the inertial relief of the wingbox mass, multiplied by the load factor of the given flight condition.

External Mass Loads model the inertial relief of externally mounted components on the wing such as leading and trailing edge devices and the engine. Each leading and trailing edge mass is coupled to one spar panel whilst the engine is coupled to both the engine mount panel and two panels on the leading edge spar. For each external mass, a resultant moment and force are computed at the centre of the coupled panel which and are then uniformly distributed over the panel's elements. This technique does not accurately represent the real distribution of loads on the coupled panels but it does apply the correct total force and moment to the wingbox structure and does not require any information about how the external mass is mounted to the wingbox. Figure 4.5 shows the location and magnitudes of the external mass loads and the resulting nodal forces on the engine mount.

Fuel Loads are applied as tractions on the lower skin which depend on the total fuel mass and the volume of each wingbox bay. For this work I use the same fuel loads as Brooks et al. [81], a full fuel load of 56 000 kg for all manoeuvre conditions and 20% of this value for the cruise condition.

Material

Throughout this thesis I use aluminium with an elastic modulus of 70 GPa, Poisson ratio of 0.3, density of 2780 kg m^{-3} and yield strength of 420 MPa for the wingbox material. Although the latest generation of transport aircraft sport almost entirely composite wing structures, composites significantly increase the complexity of the structural design parameterisation and the resulting optimisation problems. Studying only aluminium structures therefore allows this work to focus solely on differences in analysis and optimisation results caused by geometric nonlinearity rather than the minutiae of composite design optimisation.

4.1.4. Meshes

Both the FE and CFD models have coarse, medium and fine meshes. The CFD meshes contain approximately 1, 3 and 9 million cells respectively with both multiblock and overset versions. In this work I use

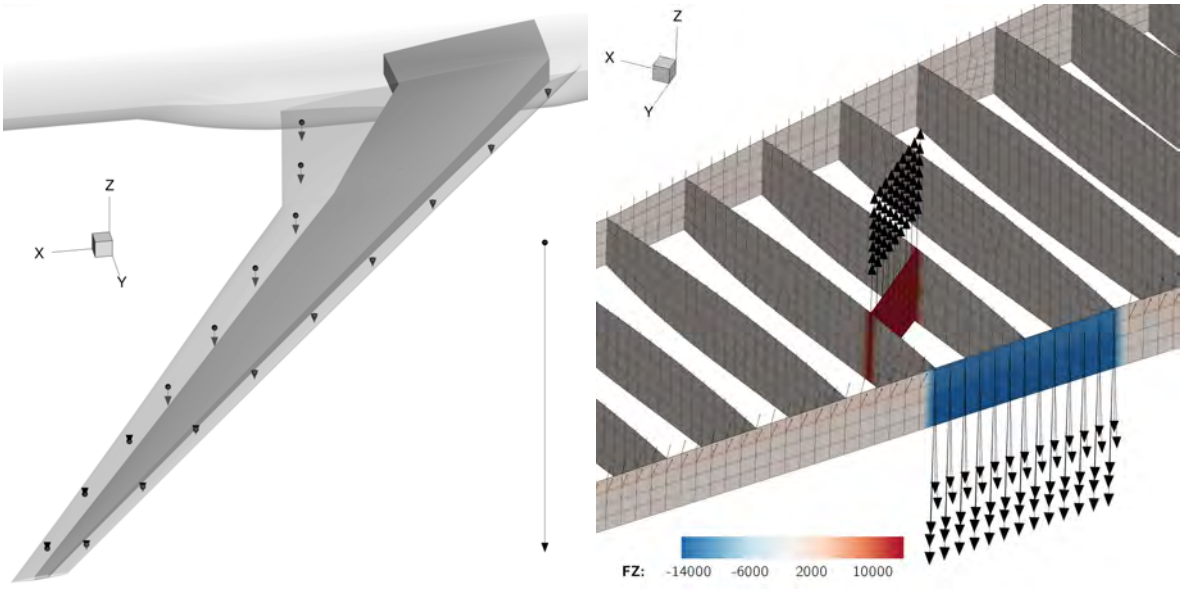

(a) Locations of external masses, including engine, for the uCRM-13.5 and inertial force vectors for the 2.5 g manoeuvre condition.

(b) Resultant nodal forces on the engine mount and leading edge spar due to engine mass during 2.5 g manoeuvre.

Figure 4.5: External mass loads on the uCRM-13.5.

the multiblock meshes.

The FE meshes consist of 4-node, second order, MITC shell elements, with the coarse, medium and fine meshes having approximately 1.5, 3.5 and 5.5×10^5 degrees of freedom.

4.2. Structural Analysis Studies

In this section I compare the results of linear and nonlinear structural analysis under fixed loading on both uCRM models. These results allow for the identification and quantification of the effects of purely structural geometric nonlinearities on the displacements and stresses in the uCRM wingboxes such that in the later aerostructural analyses, these purely structural nonlinearities can be distinguished from coupled aerostructural nonlinearities.

4.2.1. Methodology

I generate the aerodynamic loads for the structural analyses by running trimmed aerostructural analyses of the uCRM models in the three manoeuvre conditions described previously, with a linear FE formulation, and using the fine FE and CFD meshes. Figure 4.6 shows the aerodynamic traction vectors generated in the 2.5 and -1 g manoeuvre conditions, superimposed on the undeformed uCRM-13.5. Note that these traction vectors must be mapped onto the undeformed wing in order to compute the correct rigid link vectors for the load transfer calculations. The orientations of the tractions are however consistent with the deformed wing shape, thereby capturing follower force effects. The remaining loads are applied as described in the previous section.

I use the coarse wingbox meshes for these studies to reduce computational cost. Brooks et al. [81] showed that stresses were more or less constant across the three structural mesh levels.

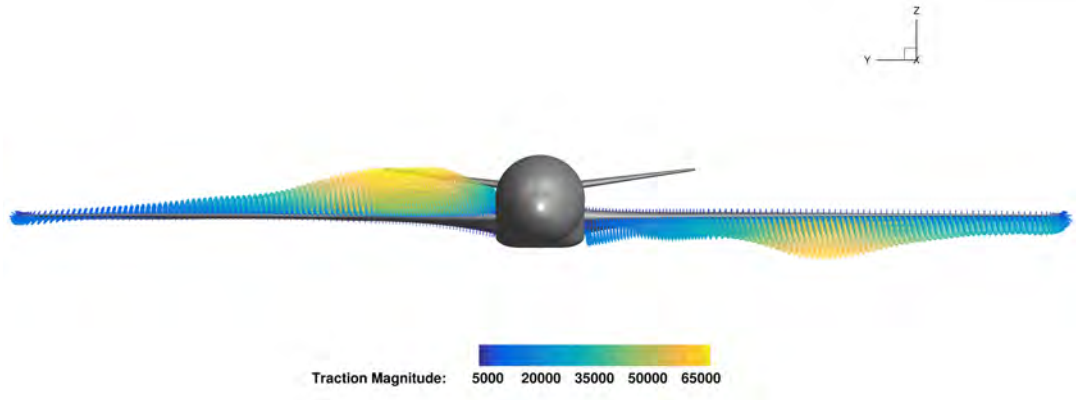

Figure 4.6: The aerodynamic loads used for structural analysis and optimisation, obtained using aerostructural analysis of the uCRM models with linear structures.

4.2.2. Results

Table 4.3 summarises the maximum failure criteria values in each wingbox component family in the 4 different structural analyses performed. The uCRM-13.5 shows some significant differences in stress and buckling criteria with a 25-40% increase in the Von Mises stresses in the skins and spars, and the buckling criteria in the ribs in the nonlinear analysis. The buckling criteria in the skins and spars do not however show a similar increase, remaining within a few percent of the values in the linear analysis. The uCRM-9 shows some similar trends, but to a lesser extent. The upper skin sees a smaller 8% increase and the spars an almost identical 24% increase in Von Mises stress while the increase in the maximum rib buckling criteria reduces to 8%.

Table 4.3: Component-wise comparison of the maximum failure criteria values from the linear and nonlinear structural analyses of the uCRM models.

		uCRM-13.5			**uCRM-9**		
Component	**Failure Criteria**[1]	Linear	Nonlinear	Δ (%)	Linear	Nonlinear	Δ (%)
Upper Skin	**Von Mises**	0.784	1.084	38.3	1.145	1.236	7.9
	Buckling	0.950	0.953	0.3	0.840	0.837	-0.4
Lower Skin	**Von Mises**	0.826	1.031	24.9	1.272	1.286	1.1
	Buckling[2]	0.944	0.918	-2.8	0.581	0.564	-3.0
Spars	**Von Mises**	0.810	1.012	25.0	0.803	0.996	24.0
	Buckling	1.183	1.143	-3.4	1.173	1.444	23.1
Ribs	**Von Mises**	0.369	0.419	13.4	0.798	0.715	-10.3
	Buckling	1.148	1.452	26.5	1.752	1.898	8.3
	Solution Time (s)	3.616	85.841	2273.9	4.497	81.947	1722.3

[1] Failure values shown are the maxmimum element averaged values from each component, normalised by either the yield stress or critical buckling stress.
[2] Lower skin buckling value taken from the -1 g condition.

These differences suggest that the nonlinear analysis produces larger bending moments in the wingbox, resulting in the increased stresses in the skins and spars, and also that it results in significant Brazier loads in the ribs, leading to the increase in their buckling loads. However, viewing the maximum stress and buckling values in each section of the wingbox gives limited insight into the real differences in the stress

field, particularly when such maxima are very sensitive to unphysical stress concentrations due to modelling simplifications or poor quality elements. In the next section I therefore compare the stress and buckling distributions over the whole wingbox.

Also included in the table are the times to compute the 2.5 g solutions on 4 cores of an Intel i7-7700HQ laptop CPU. The nonlinear solution times for the uCRM-13.5 and uCRM-9 are approximately 23 and 17 times longer than their linear counterparts respectively. This slowdown factor correlates exactly with the number of Newton-Raphson iterations required for eachnonlinear solution.

uCRM-13.5

Figure 4.7 shows the distribution of the Von Mises stress failure criterion in the uCRM-13.5 wingbox. The wingbox on the right of the figure shows the linear stress subtracted from the nonlinear stress, leaving red regions indicating areas with higher stress in the nonlinear case while blue areas show areas where stress is lower in the nonlinear case.

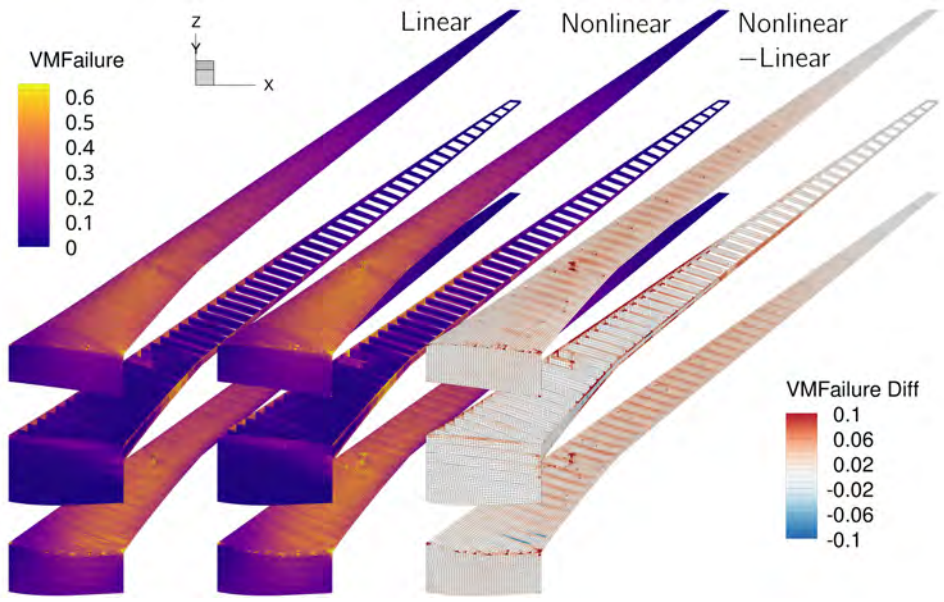

Figure 4.7: The uCRM-13.5 shows an almost uniform 10-20% increase in stress between the linear and nonlinear analyses.

The presence of almost no blue in the rightmost wingbox indicates that the stresses in the nonlinear analysis are equal to or higher than those in the linear analysis throughout the entire wingbox. In particular, the stresses over the majority of the upper and lower skins are higher by approximately 5-10% of the material yield stress, a roughly 10-20% relative increase over the linear case. There are a number of smaller areas which show a much larger increase in stress from the linear to nonlinear cases, including the sections of the skins around the engine mount and in multiple segments of the front and rear spars. These are all regions subject to external mass loads, which are computed only in the undeformed configuration and do no rotate with the structure. As such, in the nonlinear case, as the wing deforms and rotates, the previously in-plane external mass loads develop out-of plane components, resulting in significantly greater stresses. This is an unrealistic phenomena but developing a geometrically exact method for applying these external loads is considered outside the scope of this work.

Figure 4.8 compares the value of the buckling failure criterion from the same analyses. The buckling criterion values in the upper skin are approximately 10-20% higher in the nonlinear case, corresponding directly to the increase in compressive stress. More interestingly, a large number of the ribs in the midspan of the wing show a relative increase in the buckling criterion of around 50%, strongly suggesting the presence of Brazier loading.

To confirm this suggestion, figure 4.9 compares the axial stress in the stiffener-wise direction of the ribs.

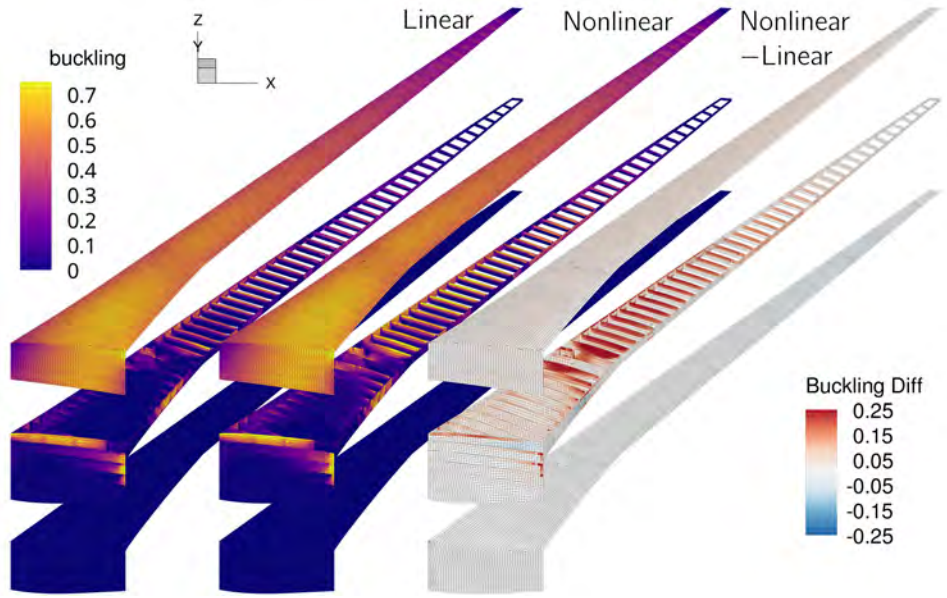

Figure 4.8: The uCRM-13.5 ribs show around a 50% increase in buckling criteria in the nonlinear analysis, suggesting significant Brazier loading.

The stress shown is computed at the mid-section of the shell, thus negating any axial stresses due to bending. The results show a drastic increase in the compressive stress in the ribs, particularly in the region just outboard of the engine mount where the buckling criteria also increased greatly. The compressive stresses in the linear analysis, if present at all, are smaller than in the nonlinear case by at least an order of magnitude.

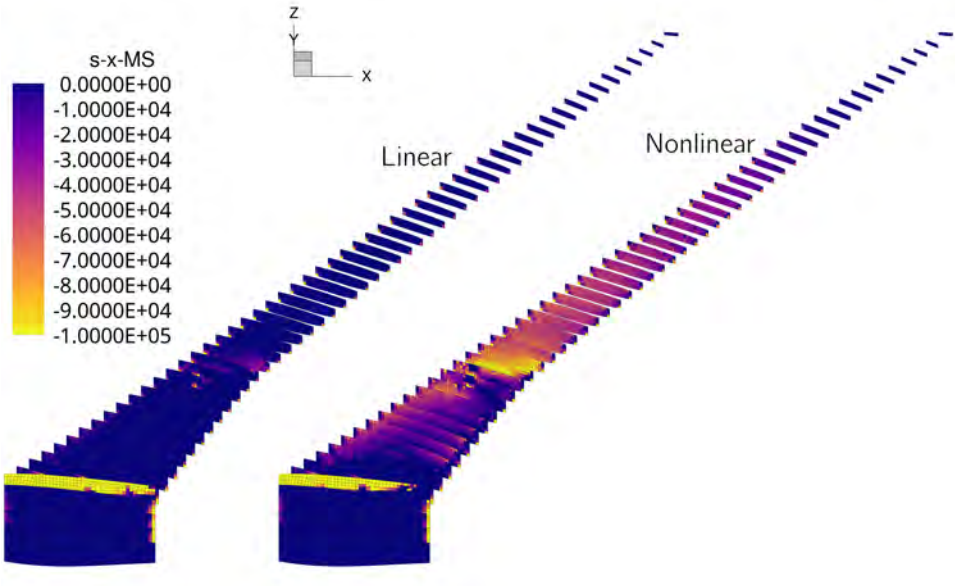

Figure 4.9: The presence of Brazier loads is confirmed by examining the axial stress in the ribs.

Comparing the deformations of the two wings, shown in figure 4.10 provides some clues into the cause of the higher stresses seen in the nonlinear analysis. The nonlinear deformed shape shows the expected tip-shortening effect but also appears to show a higher curvature, particularly in the outboard portion of the wing. This greater curvature suggests a higher bending moment within the wing which is consistent with the follower force phenomena described in section 1.1.3 wherein the rotation of the lift forces on the outboard portion of the wing incorrectly result in a decrease in the bending moment seen by linear

structural models due to their consideration of only the undeformed configuration of the wingbox. This increased bending moment is then responsible for the relatively uniform increase in stress seen in the upper and lower skins of the wingbox.

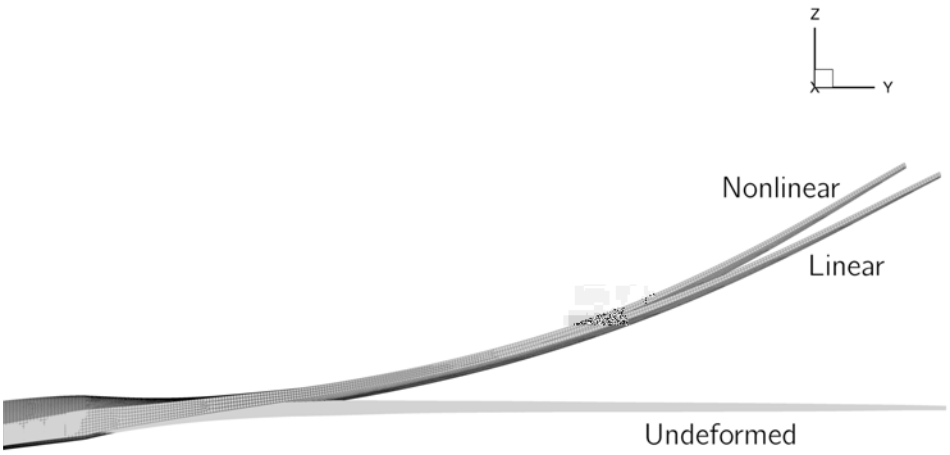

Figure 4.10: The nonlinear analysis shows the expected tip-shortening phenomenon and shows greater curvature, indicating a greater bending moment in the wing.

Figure 4.11 shows the chordwise, spanwise and vertical components of the displacement and the wing twist, all extracted from points along a line on the upper skin of the wingbox. The twist distributions in the linear and nonlinear case are almost identical, with nonlinear case showing slightly more washout. This suggests that the higher bending curvature in the nonlinear case results in more washout due to the wing's geometric bend-twist coupling and is more than enough to cancel out any decrease in washout due to drag-torsion effects. Also shown in figure 4.11 are the chordwise (ΔX), spanwise (ΔY) and vertical (ΔZ) components of the displacement. From these curves it is possible to quantify the effective tip-shortening in the nonlinear case, around 4% of the wing's semispan. Also of note is that the chordwise deformation is reversed between the two cases with the wing tip moving forwards in the nonlinear case. This is likely due mostly to the same bending kinematics that cause the tip shortening effect of the wing, which result in a chordwise displacement due the wing's sweep.

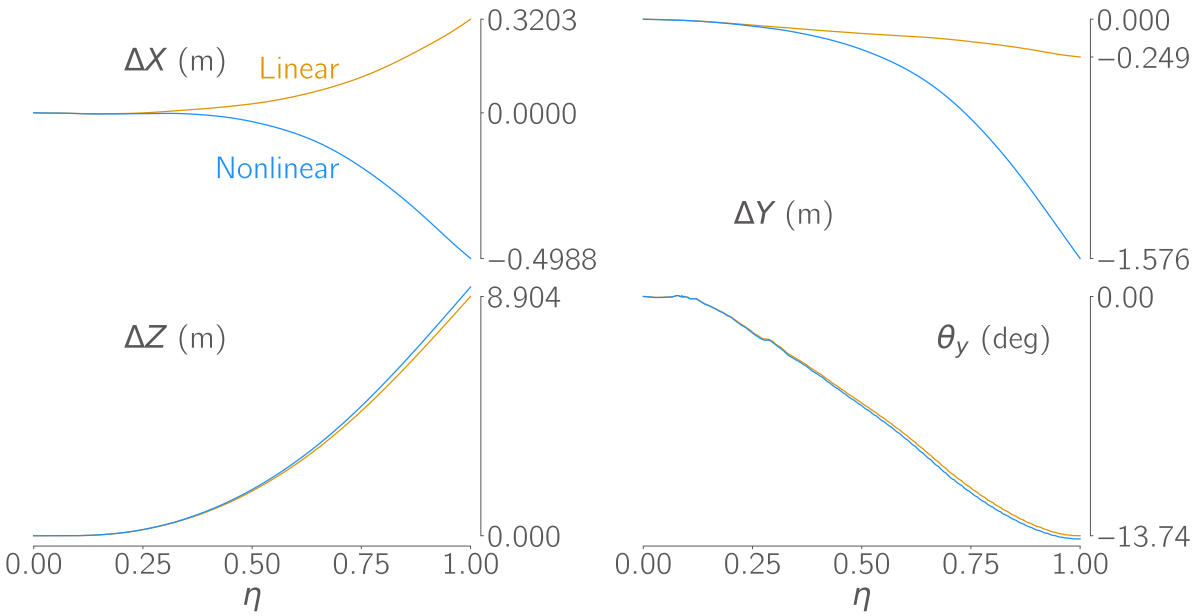

Figure 4.11: Spanwise deformation distributions. The nonlinear wing shows slightly more washout

uCRM-9

Figures 4.12-4.15 show the same comparisons as above between the linear and nonlinear structural analyses performed on the uCRM-9. The differences in stresses and deformations in the lower aspect-ratio wingbox of the uCRM-9 are similar to those seen in the uCRM-13.5 but with a lower magnitude. In the nonlinear case, the stresses in skins are around 5% higher and the buckling criterion is again approximately 50% higher in ribs. The nonlinear analysis also again shows slightly more washout but by an essentially negligible margin and the tip shortening reduces to around 2% of the wing semispan.

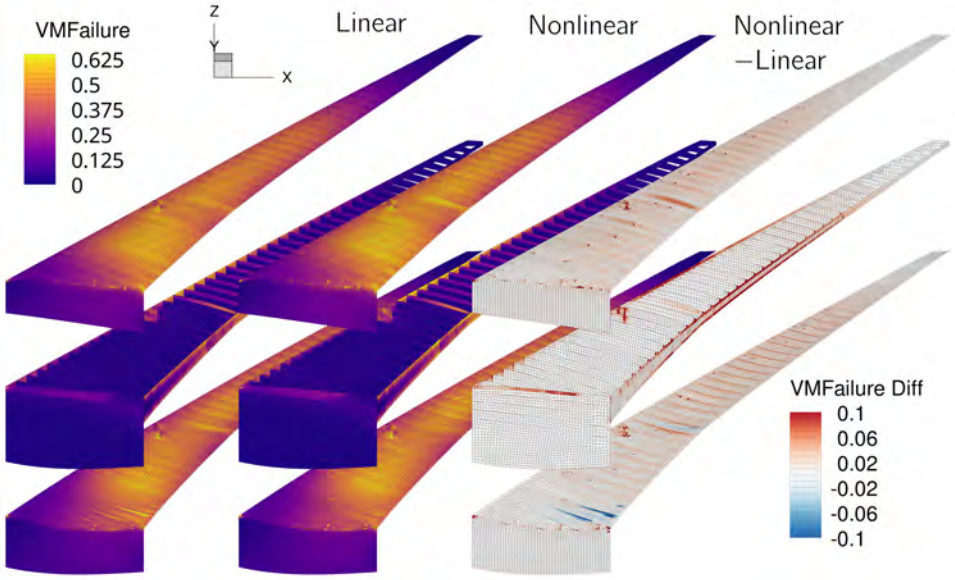

Figure 4.12: The uCRM-9 shows an almost uniform 5-10% increase in stress between the linear and nonlinear analyses.

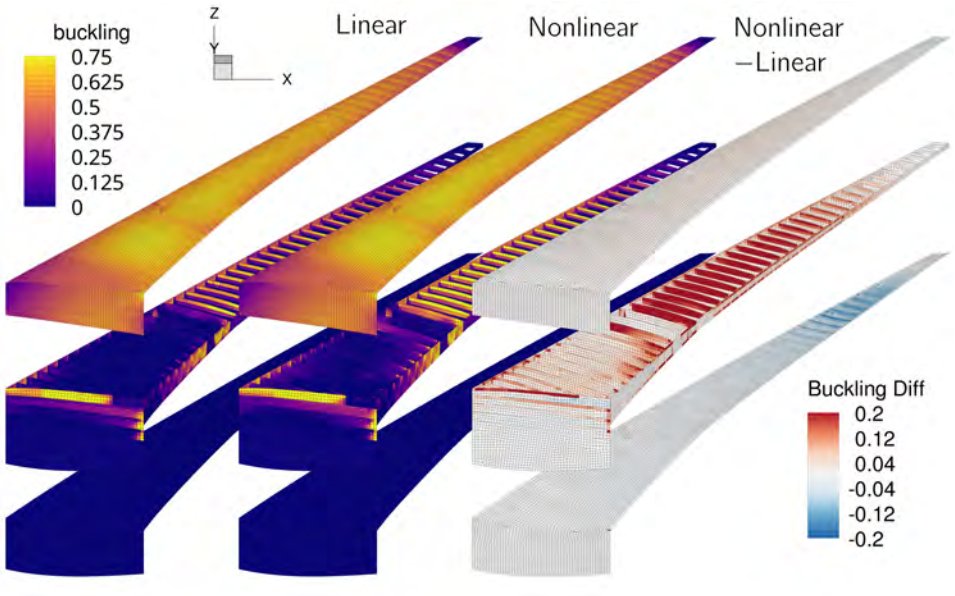

Figure 4.13: The uCRM-9 ribs also show around a 50% increase in buckling criteria in the nonlinear analysis.

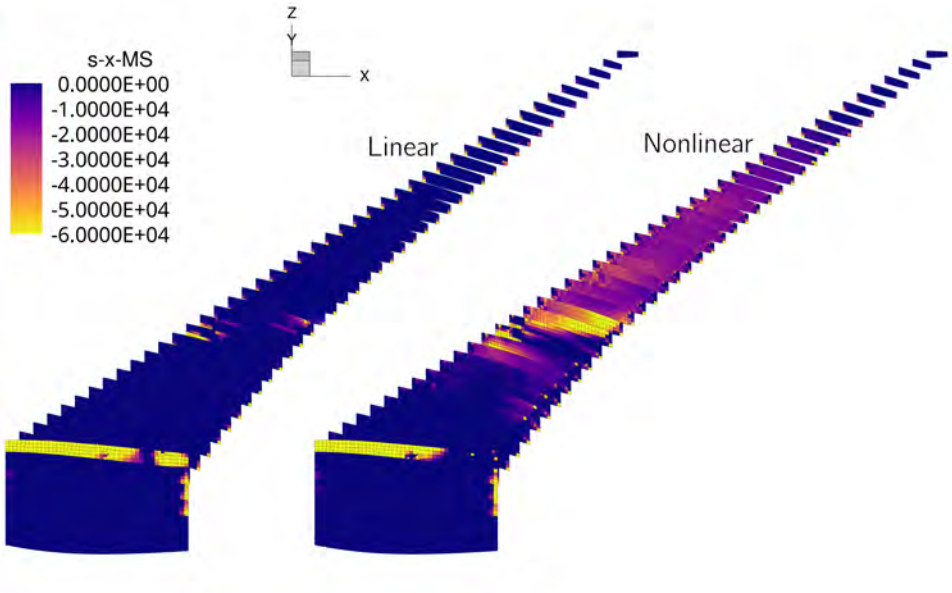

Figure 4.14: Again, the presence of Brazier loads is confirmed by examining the axial stress in the uCRM-9 ribs.

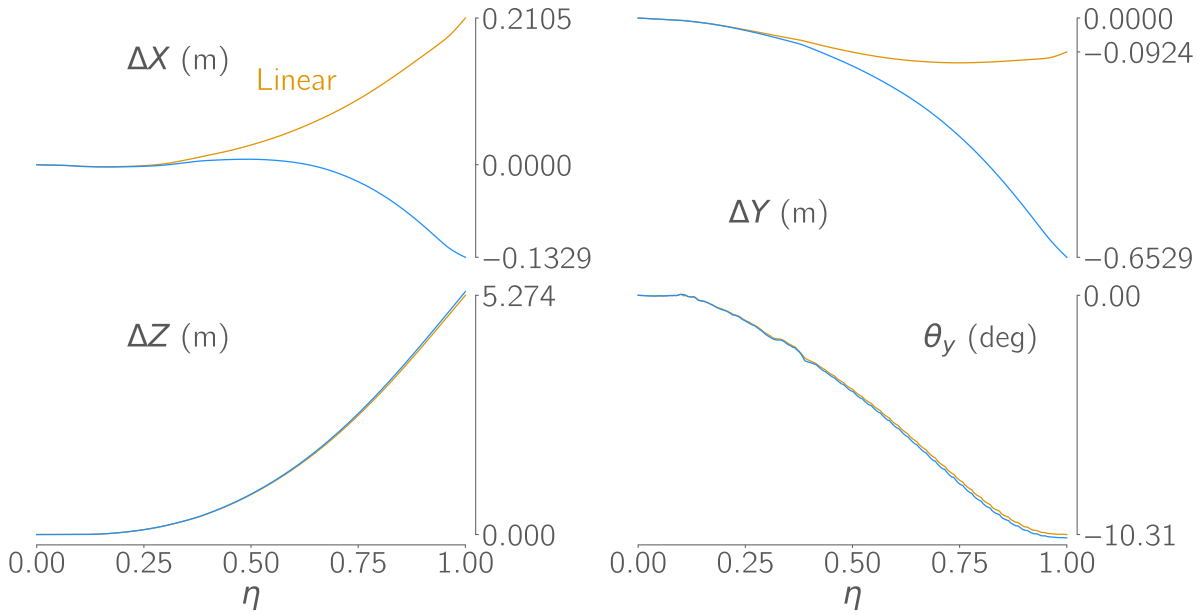

Figure 4.15: Spanwise deformation distributions from the uCRM-9.

4.3. Aerostructural Analysis Studies

4.3.1. Methodology

Untrimmed Analysis

I perform aerostructural analysis of both uCRM models with both linear and nonlinear structural formulations. In this first study, I analyse both the linear and nonlinear models at the same angle of attack, which produces the correct 2.5 g load factor with the linear model. I use the fine CFD mesh to maximise the accuracy with which the aerodynamic forces are computed but retain the coarse FE mesh.

Trimmed Analysis

In this study, I repeat the aerostructural analyses above, now trimmed such that the linear and nonlinear models achieve the same lift coefficient and a zero pitching moment coefficient. To this point, the studies performed have focused only on the effect of geometric nonlinearities in manoeuvre situations. Although the lower load factor of cruise flight means that any differences in deformed wing shapes should be reduced, drag can be very sensitive to small geometry changes. I therefore also perform a trimmed analysis at the cruise condition shown in table 4.2 to assess the necessity of using geometrically nonlinear analysis for accurate cruise drag prediction. I again use the same combination of fine CFD and coarse FE meshes. In this case the fine CFD mesh is particularly important for accurate cruise drag prediction.

To trim the aircraft fully, the rotation of the horizontal tail is parameterised using the free-form deformation (FFD) approach described in section 2.6. Figure 4.16 shows the FFD volumes used to parameterise the full uCRM-13.5 geometry. To rotate the tail, the highlighted FFD volume surrounding the tail is rotated rigidly around the global Y axis.

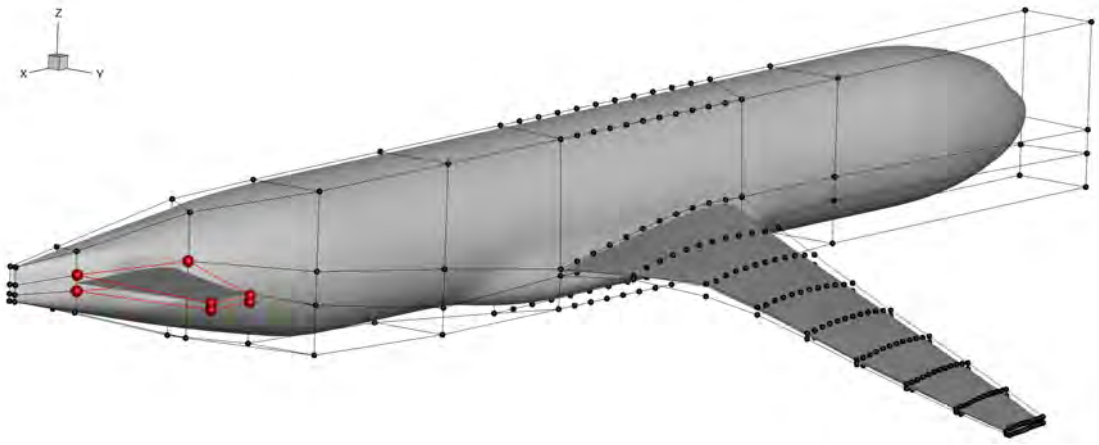

Figure 4.16: The FFD volumes used to parameterise the uCRM-13.5, with the volume controlling tail rotation highlighted.

Each aerostructural analysis is judged to have converged when the residual norms of both the CFD and FE problems has reduced by a factor of 10^{-8}. The trimmed solver is judged to have converged when both the lift and moment coefficients are within 10^{-4} of their respective targets. I run the nonlinear aerostructural analyses with a 5 iteration load ramp and use a conservative initial under-relaxation factor of 0.1 for all analyses in order to avoid the formation of negative volume CFD cells in the early Gauss-Seidel iterations.

4.3.2. Results

Table 4.4 summarises the lift, drag and pitching moment coefficients (C_L, C_D and C_{M_y}) from the four untrimmed analyses along with the coordinates of the centre of pressure (COP) of the wing relative to the uCRM's centre of gravity (COG)[1]. Each of the aerodynamic coefficients are also broken down into the individual contributions from the wing, fuselage and horizontal tail. Also included are the chordwise spanwise and vertical wingbox tip displacements, demonstrating the span shortening effect in the nonlinear analyses.

Untrimmed Analysis

In both uCRM models, span shortening has a small but noticeable effect on a number of the aerodynamic quantities. The lift coefficient for both models is 1.7-1.8% lower, with the vast majority of the decrease in lift coming from the wing. These decreases in lift are slightly smaller than the decrease in effective span of

[1]The COP is not a unique point and can be defined anywhere along the line of the resultant aerodynamic force. In this work I compute the location of the COP in the same horizontal plane as the COG

Table 4.4: Aerodynamic functions of interest from the untrimmed aerostructural analysis studies

	uCRM-13.5			**uCRM-9**			
Value	Linear	Nonlinear	Δ (%)	Linear	Nonlinear	Δ (%)	**Units**
C_L	**0.595**	**0.585**	**-1.7**	**0.629**	**0.618**	**-1.8**	
$C_{L,wing}$	0.519	0.510	-1.8	0.545	0.535	-2.0	
$C_{L,fuse}$	0.087	0.087	0.4	0.092	0.092	-0.5	
$C_{L,tail}$	-0.011	-0.012	8.1	-0.009	-0.009	0.2	
C_D	**298.4**	**301.6**	**1.1**	**323.3**	**324.2**	**0.3**	**cts**
$C_{D,wing}$	137.5	141.0	2.5	158.7	159.5	0.5	cts
$C_{D,fuse}$	152.8	152.1	-0.5	155.7	156.1	0.3	cts
$C_{D,tail}$	4.2	2.9	-30.1	4.5	5.0	10.4	cts
C_{M_y}	**0.173**	**0.204**	**17.8**	**0.040**	**0.056**	**40.6**	
$C_{M_y,wing}$	-0.036	-0.010	-70.9	-0.111	-0.095	-14.4	
$C_{M_y,fuse}$	0.160	0.160	0.3	0.118	0.118	0.1	
$C_{M_y,tail}$	0.049	0.054	8.9	0.033	0.033	0.4	
$COP_{wing,x}$	**0.402**	**0.117**	**-71**	**1.454**	**1.270**	**-12.6**	**m**
$COP_{wing,y}$	**13.78**	**13.35**	**-3.1**	**12.89**	**12.65**	**-1.9**	**m**
ΔX_{tip}	0.41	-0.49	-220.5	0.28	-0.10	-135.7	**m**
ΔY_{tip}	-0.23	-1.60	593.1	-0.02	-0.64	2591.5	**m**
ΔZ_{tip}	10.39	9.61	-7.5	6.00	5.70	-5.1	**m**

the wing, this is likely due to the fact that most of the span shortening occurs in the outboard sections of the wing whilst most of it's lift is generated on the inboard portion. As explained in the previous section, because of the backwards sweep of the uCRM wing, the span shortening in the nonlinear cases shifts the centre of pressure not only inwards but also forwards. This leads to significant increases in the total nose up pitching moment of the aircraft, with 18 and 40% increases in the uCRM-13.5 and 9 respectively. It should be noted, however, that the 40% change in the uCRM-9 pitching moment represents a smaller absolute change than seen with the uCRM-13.5.

The spanwise lift distributions, when plotted in absolute coordinates as in figure 4.17, are almost indistinguishable on the inboard half of the span whilst the nonlinear cases show lower lift on the outboard half of the wing. It is unclear from these plots, however, whether this decrease in lift at a given spanwise station is purely a result of the nonlinear span shortening or if it is due to the nonlinear analysis exhibiting more passive load alleviation. Plotting the normalised lift distributions against the normalised spanwise coordinate in figure 4.18 removes differences in the curves due to different spans and different total lift levels, allowing a more straightforward comparison of the relative spanwise distribution of lift. With these differences accounted for, the linear and nonlinear lift distributions become even less distinguishable. The slight differences in the uCRM-13.5 curves show that the linear case produces a greater normalised lift over the inboard 30% whilst the nonlinear wing produces a greater a greater normalised lift over the 30-75% span section. This suggests that the nonlinear case sees less passive load alleviation, although the difference appears to be almost negligible. Another observation in support of this theory is that the inboard shift

in the COP in the nonlinear case (3.1%) is slightly less than the relative reduction in wing span (3.8%).

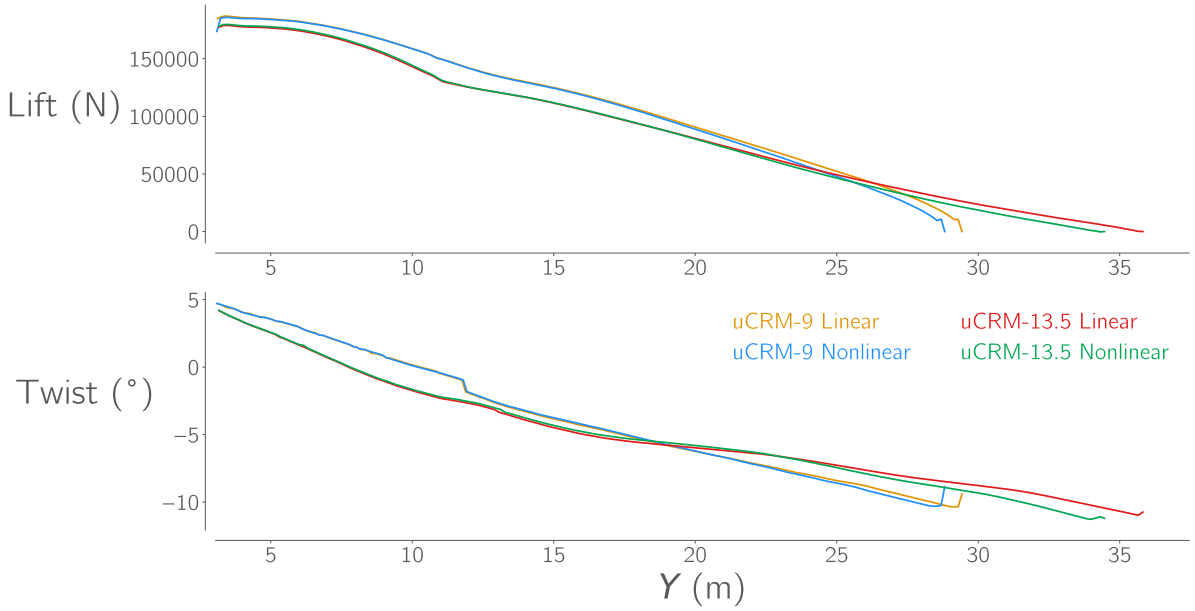

Figure 4.17: Comparison of the lift and twist distributions from the untrimmed aerostructural analyses, plotted against the absolute spanwise coordinate.

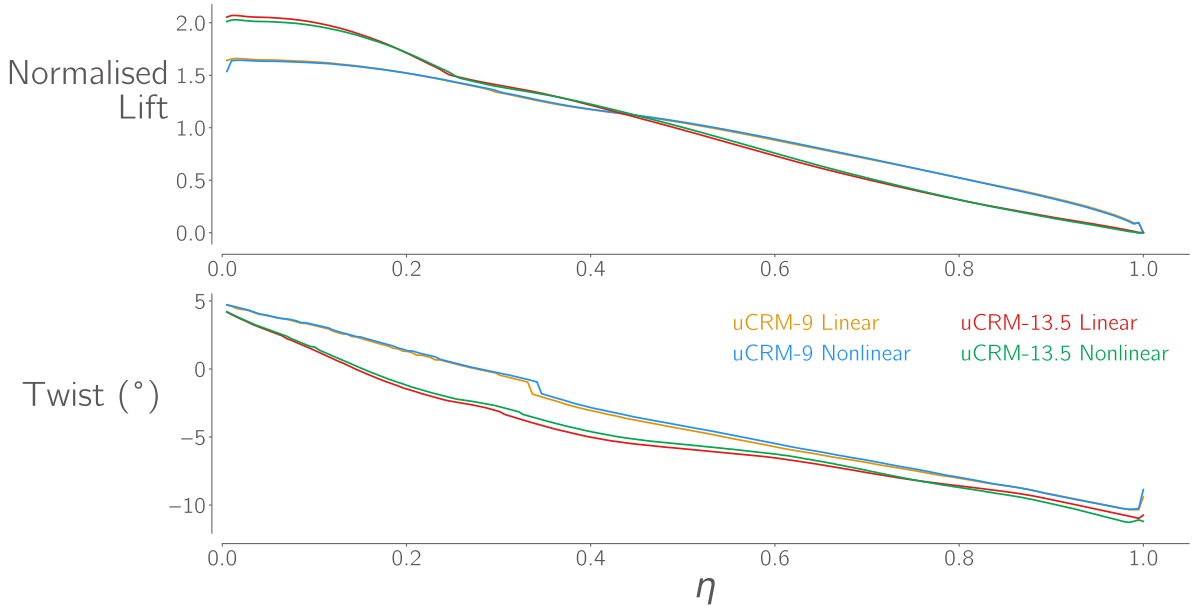

Figure 4.18: Normalising the lift and spanwise coordinate shows that there is very little difference in the relative spanwise lift distributions in the linear and nonlinear analyses.

Table 4.5 shows the number of Gauss-Seidel iterations required for each solution and the solution times, including the breakdown of the time spent in the 3 most time consuming portions of the solver, the CFD analysis, FE analysis and the mesh warping. Of the 4 untrimmed analyses run, only the nonlinear analysis on the uCRM-9 achieved the specified convergence tolerance, while the other 3 analyses all took the maximum allowable number of Gauss-Seidel iterations. Both linear cases ran in similar times with almost all of the difference in time spent in the CFD solver. The time spent on mesh warping and FE analysis is almost identical for the two case as, once the linear stiffness matrix has been factorised, both are fast explicit calculations. The nonlinear analysis of the uCRM-13.5 takes just under 40% longer than the linear analysis. 73% of this additional time is spent in the nonlinear FE analysis, a 14-fold increase in

the total FE analysis time compared to the linear case. Although significant, this increase is smaller than the approximately 20-fold increase seen in the structural analysis studies which indicates that, for at least some Gauss-Seidel iterations, the nonlinear structural solver is able to save time by restarting from the previous structural solution.

The two uCRM-9 analyses take approximately the same time to complete. Correcting for the lower number of iterations required for the nonlinear solution however gives the same time increase per Gauss-Seidel iteration for the nonlinear solution, around 40%.

It should be noted that, due to the use of the fine CFD meshes, these analyses were run on 192 CPU cores. This is well beyond the optimum number of cores for the structural analysis due to the significantly smaller number of elements in the structural mesh. As a result, it is possible that the solution time could be decreased by running the structural analyses a smaller subset of the cores used by the aerostructural solver.

Table 4.5: Solution times from the untrimmed aerostructural analysis studies

	uCRM-13.5			**uCRM-9**			
Value	Linear	Nonlinear	Δ (%)	Linear	Nonlinear	Δ (%)	**Units**
GS iterations	60	60	0.0	60	36	-40.0	
Total Solve time	1184	1638	38.3	1112	1134	1.9	s
CFD Time	933	1059	13.6	854	760	-10.9	s
FE Time	22	343	1443.4	23	231	912.2	s
Warp Time	22	22	0.2	23	13	-41.1	s

Trimmed Analysis

Table 4.6 summarises the same quantities of interest as table 4.4 for the trimmed 2.5 g manoeuvre analyses, with the addition of the angle of attack and tail rotation angle required to trim each case. Consistent with the reduction in lift due to span shortening seen in the untrimmed analyses, the trimmed nonlinear analyses require higher angles of attack to achieve the target lift coefficient, although the differences are on the order of 1%. There is, again, a noticeable shifting forward and inboard of the COP which causes most of the largest relative changes to be in quantities related to the pitching moment of the aircraft. For example, in both models, the nonlinear cases produce around 1% less lift from the wing than the linear cases. This is partially due to a small increase in the lift generated by the fuselage at the higher angle of attack and partially because the forward shift in the COP of the nonlinear wings necessitates more lift, or less downforce, from the horizontal tail. In the uCRM-13.5, this results in around half of the 8 count drag increase coming from the tail.

As with the untrimmed results, however, what appear as large relative changes in some values are, in reality, fairly small absolute changes. For example, the 55% increase in the tail angle required to trim the nonlinear uCRM-13.5 case amounts to only 0.3° rotation.

I omit any comparison of the lift and twist distributions from the trimmed analyses as they are largely indistinguishable from those from the untrimmed analyses shown in figures 4.17 and 4.18. This is unsurprising given the relatively small changes in angle of attack and tail rotation required to trim the uCRM models.

Table 4.6: Aerodynamic functions of interest from the trimmed aerostructural 2.5 g manoeuvre analyses

	uCRM-13.5			**uCRM-9**			
Value	Linear	Nonlinear	Δ (%)	Linear	Nonlinear	Δ (%)	**Units**
α	**5.545**	**5.582**	**0.7**	**5.527**	**5.608**	**1.5**	**°**
θ_{tail}	**0.554**	**0.863**	**55.7**	**-0.872**	**-0.678**	**-22.2**	**°**
C_L	**0.587**	**0.587**	**0**	**0.587**	**0.587**	**0**	
$C_{L,wing}$	0.487	0.482	-1.1	0.506	0.501	-0.8	
$C_{L,fuse}$	0.083	0.084	1.6	0.084	0.085	1.2	
$C_{L,tail}$	0.017	0.021	24.3	-0.003	0.000	-105.7	
C_D	**277.8**	**285.5**	**2.8**	**281.5**	**288.1**	**2.4**	**cts**
$C_{D,wing}$	113.0	115.5	2.3	131.9	134.2	1.7	cts
$C_{D,fuse}$	136.8	138.2	1.1	136.9	138.7	1.3	cts
$C_{D,tail}$	24.3	28.0	15.2	9.6	11.8	22.8	cts
C_{M_y}	**0**	**0**	**-**	**0**	**0**	**-**	
$C_{M_y,wing}$	-0.041	-0.020	-50.7	-0.110	-0.097	-11.4	
$C_{M_y,fuse}$	0.124	0.123	-1.2	0.100	0.100	0.1	
$C_{M_y,tail}$	-0.083	-0.103	23.3	0.010	-0.003	-131.2	
$COP_{wing,x}$	**0.490**	**0.244**	**-50.2**	**1.548**	**1.383**	**-10.7**	**m**
$COP_{wing,y}$	**13.758**	**13.398**	**-2.6**	**12.893**	**12.684**	**-1.6**	**m**
ΔX_{tip}	**0.384**	**-0.419**	**-209.2**	**0.268**	**-0.071**	**-126.6**	**m**
ΔY_{tip}	**-0.217**	**-1.415**	**551.9**	**-0.027**	**-0.562**	**1981.9**	**m**
ΔZ_{tip}	**9.614**	**9.027**	**-6.1**	**5.544**	**5.344**	**-3.6**	**m**

Figures 4.19 and 4.20 show the same comparisons of the Von Mises and buckling failure criteria in the uCRM-13.5 wingbox as were shown in the structural analysis studies, now for the trimmed 2.5 g manoeuvre cases. In stark contrast to the results from the structural analysis studies, the Von Mises stresses resulting from the nonlinear aerostructural analyses are lower than for the linear analysis. As was the case in the structural analysis studies, the difference in Von Mises stress is relatively uniform over the upper and lower skins of the wingbox and of a similar magnitude, varying from 2-5% of the material yield stress, representing a relative decrease of 4-10%. The stresses in the ribs and spars remain higher in the nonlinear case, as do the stresses around the engine mount which, as explained during the structural analysis studies is due to inaccuracies in the way that the external mass loads are modelled in these analyses.

The differences in the buckling criteria between linear and nonlinear analyses are more similar to those seen in the structural analysis studies. The decrease in buckling loads in the upper skin of the wingbox is on the order of 1% in the buckling critical areas near the wing root. Although this difference grows towards the wing tip, these areas are not buckling critical. As was the case in the structural analysis studies, the buckling loads in the ribs are again approximately 50% higher in the nonlinear case due to Brazier loads.

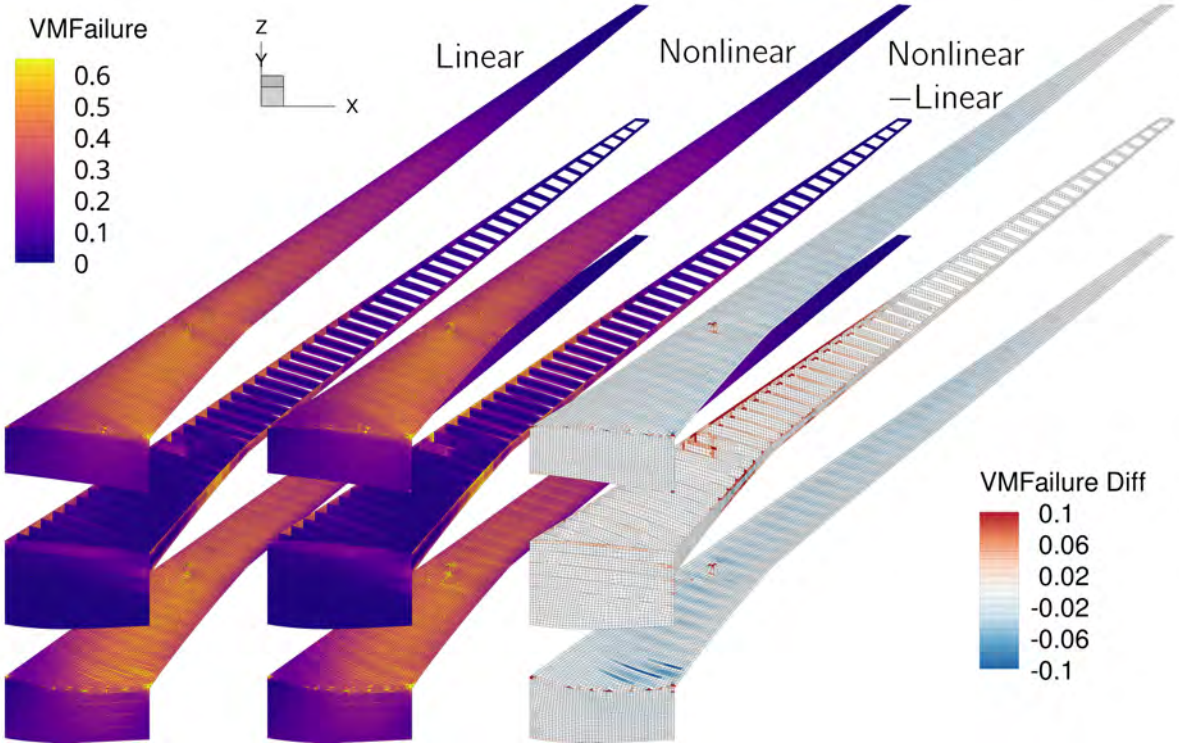

Figure 4.19: The nonlinear aerostructural analysis results in a significant reduction in the wingbox bending moment and resulting stresses, in stark contrast to the results of the pure structural analyses.

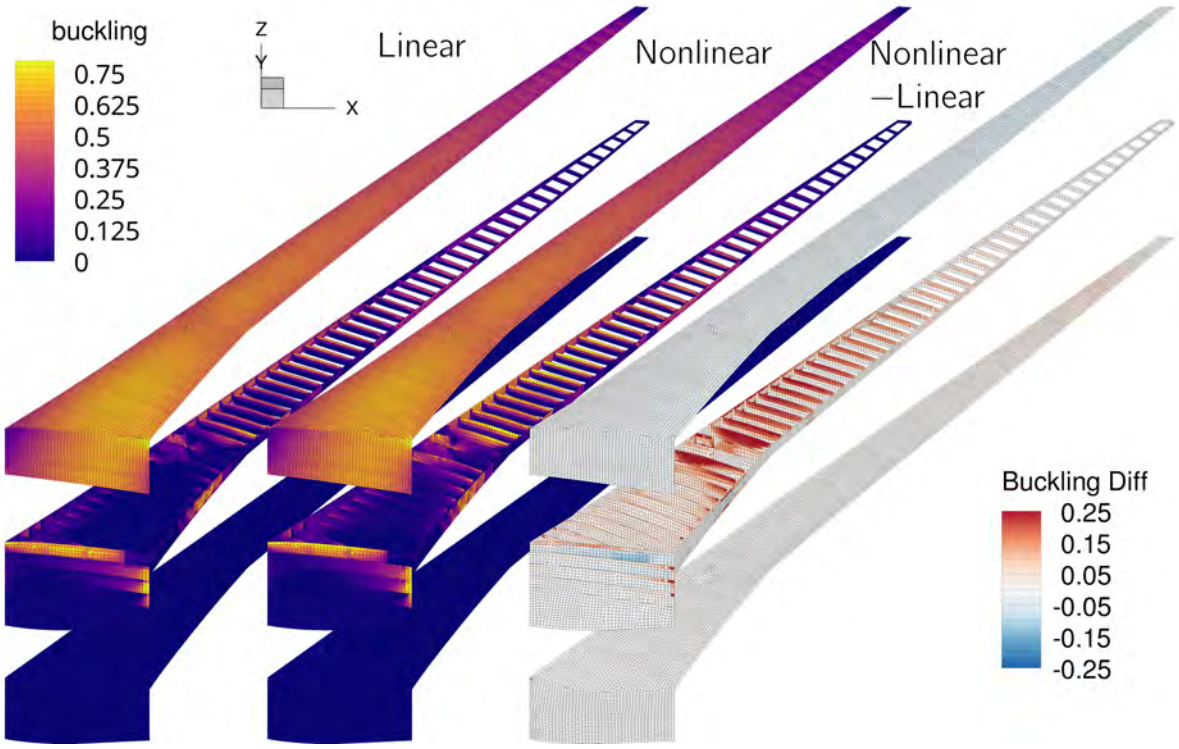

Figure 4.20: The buckling loads in the upper wing skin are more similar between the linear and nonlinear cases than the Von Mises stress. The ribs again see a 50% increase in buckling loads in the nonlinear case due to Brazier loads.

The most obvious explanation for the stress decreases seen in the nonlinear analysis is the inboard shift

in COP induced by the nonlinear analysis's span shortening. Combining the relative reductions in the lift produced by the wing and the spanwise coordinate of the COP of the uCRM-13.5 gives a rough approximation of the relative decrease in the lift induced bending moment in the wingbox of 3.7%. Accounting for the reduced lift force moment arm due to the constraint on the wingbox at the fuselage junction increase this value to 4.6%, which correlates with the lower end of the range of observed stress decreases in the nonlinear analysis. This explanation also goes some way to accounting for the 6% decrease in the wingtip vertical deflection seen in the nonlinear case. An lower level explanation of the same phenomena is that, in the linear analyses, the linearised bending kinematics cause the outboard portions of the wing to be artificially stretched, this increase in area then results in increased forces in the outboard section of the wing when integrating the aerodynamic pressure over the wing's surface, causing the observed increase in the wing bending moment.

Finally, in table 4.7 I compare the results of the trimmed analyses performed at the cruise flight condition, the purpose of which is to assess the effect of geometric nonlinearity on cruise drag prediction. Perhaps unsurprisingly, the differences between the linear and nonlinear results are similar to those seen in the manoeuvre condition analyses but significantly smaller. For both uCRM models, between linear and nonlinear analyses, the difference in drag coefficient is below 1 drag count, in tail rotation below 0.1°, and in angle of attack below 0.02°.

Table 4.7: Functions of interest from the trimmed aerostructural cruise analyses

	uCRM-13.5			**uCRM-9**			
Value	Linear	Nonlinear	Δ (%)	Linear	Nonlinear	Δ (%)	**Units**
α	**2.197**	**2.202**	**0.2**	**2.194**	**2.208**	**0.6**	**°**
θ_{tail}	**-0.036**	**0.094**	**-364.1**	**-0.788**	**-0.736**	**-6.6**	**°**
C_L	**0.500**	**0.500**	**0**	**0.500**	**0.500**	**0**	
$C_{L,wing}$	0.462	0.460	-0.6	0.472	0.471	-0.2	
$C_{L,fuse}$	0.058	0.058	0.9	0.060	0.060	0.5	
$C_{L,tail}$	-0.020	-0.018	-10.1	-0.032	-0.031	-2.7	
C_D	**216.2**	**217.1**	**0.4**	**244.9**	**245.5**	**0.2**	**cts**
$C_{D,wing}$	113.6	113.8	0.2	148.8	149.1	0.2	cts
$C_{D,fuse}$	95.5	95.5	0.1	92.8	93.0	0.2	cts
$C_{D,tail}$	4.9	5.4	10.9	1.9	2.0	4.7	cts
C_{M_y}	**0**	**0**	**-**	**0**	**0**	**-**	
$C_{M_y,wing}$	-0.160	-0.150	-6.7	-0.174	-0.171	-2.1	
$C_{M_y,fuse}$	0.066	0.065	-1.5	0.051	0.050	-0.3	
$C_{M_y,tail}$	0.094	0.084	-10.3	0.123	0.120	-2.8	
$COP_{wing,x}$	**2.000**	**1.876**	**-6.2**	**2.602**	**2.555**	**-1.8**	**m**
$COP_{wing,y}$	**15.51**	**15.33**	**-1.2**	**14.11**	**14.05**	**-0.5**	**m**
ΔX_{tip}	**0.232**	**-0.088**	**-138.1**	**0.115**	**0.034**	**-70.5**	**m**
ΔY_{tip}	**-0.125**	**-0.603**	**383.6**	**-0.001**	**-0.132**	**20357.9**	**m**
ΔZ_{tip}	**5.651**	**5.497**	**-2.7**	**2.627**	**2.598**	**-1.1**	**m**

4.4. Chapter Summary

In this chapter I presented the results multiple studies investigating the effects of including geometric nonlinearity in structural and aerostructural analysis of modern transport aircraft with moderate and high aspect-ratio wings using the undeflected common research model (uCRM)s. I posited that the uCRMs represent excellent test cases for such studies as they both the aspect-ratios typical of current day transport aircraft and those entering service in the next two decades.

In the first study I performed linear and nonlinear structural analysis of the two uCRM wingboxes using a fixed set of loads generated from the same set of geometrically linear aerostructural analyses. I showed that, in the high aspect-ratio uCRM-13.5 wingbox, the nonlinear analysis results in greater bending stresses throughout, particularly in the upper and lower skins where increases of around 10% are observed in the Von Mises stress. Greater stress increases are seen locally around areas subjected to loads from external masses which, I propose, are due to the fact that such loads are not applied in a geometrically exact manner in these analyses. I demonstrated the presence of Brazier loads in the nonlinear analyses which are evident from a substantial increase in the compressive axial stress and buckling failure criteria in the wingbox ribs. Comparing spanwise deflection distributions, I showed that the correct modelling of large bending kinematics in the geometrically nonlinear analysis results in a significant inward and forward deflection of the wingtip. I also showed that any geometrically nonlinear drag-torsion effects, if present, have little effect on the twisting deformation of the wing, which is dominated by the geometric bend-twist coupling caused by wing sweep. I proposed that the observed stress bending stress increases and the slightly increased washout in the nonlinear case are due to the nonlinear analysis correctly maintaining the bending moment caused by the lift forces on the outboard portions of the wing which are rotated inboard. This conclusion is consistent with previous works demonstrating that linear analyses underpredict bending deformation when subject to follower forces.

I also showed that similar differences in stress distributions are seen when performing nonlinear structural analysis on the moderate aspect-ratio uCRM-9 wing, albeit to a lesser extent. In this case, the increase in bending stresses is roughly halved, whilst the increase in the rib buckling criteria remains around 50%.

I then presented the results of a series of coupled aerostructural analyses of the same models, again using geometrically linear and nonlinear structural analyses. Performing untrimmed aerostructural analyses at the same angle of attack, I demonstrated that when using nonlinear structural analysis, the span shortening behaviour resulting from correctly modelled bending kinematics shifts the centre of pressure (COP) of the wing inboard and forwards, resulting in a 20-40% change in the untrimmed pitching moment of the aircraft. This span shortening also decreases the total lift produced by the wing by almost 2%. Comparing normalised lift distributions between the linear and nonlinear analyses, I showed that there is almost no difference in the amount of passive load alleviation achieved in the linear and nonlinear analyses. Comparing solution times, I demonstrated that the inclusion of nonlinear structural behaviour increases the computational cost of each coupled aerostructural analysis iteration by 40%.. In contrast, the solution slowdown of nonlinear structural analysis is almost exactly equal to the number of Newton-Raphson iterations taken to converge the analysis which, for the presented analyses is on the order of 20×.

I then performed aerostructural analyses in the same 2.5 g manoeuvre condition, with all analyses trimmed to the same lift coefficient and zero pitching moment. I demonstrated again that span shortening is the main cause of difference between the linear and nonlinear analysis results, with the nonlinear analyses requiring a small increase in angle of attack to achieve the same lift coefficient. This angle of attack was however mitigated by the need for less downforce from the horizontal tail to trim the aircraft due to the forward shift in the wing's COP, as a result, the lift produced by the wing in the nonlinear cases was still approximately 1% smaller than in the linear cases.

I then compared the wingbox stress distributions from the trimmed aerostructural analyses and showed that the nonlinear analysis resulted in 5-10% lower bending stresses in the wing skins, the opposite of the results seen in the structural analysis studies. To a first order approximation I proposed that this reduction in bending moment in the wingbox is due to a combination of the small reduction in lift produced by the wing and the inboard shift in the wing's COP. Buckling loads in the wing skins showed a similar decrease whilst the ribs saw a similar increase in buckling loads as seen in the structural analysis studies.

Finally, I ran similar trimmed aerostructural analyses in a transonic cruise condition.. The differences in key values such as the drag coefficient and trim angle of attack differed by 0.5% or less, indicating that aerostructural analysis with linear structural models is sufficient for accurate cruise drag prediction.

Looking forward, to the following optimisation, these results suggest that:

In structural optimisation, wingboxes optimised using nonlinear analysis will be heavier due to the higher stresses observed in the nonlinear structural analyses under fixed loading.

In aeroelastic optimisation, wingboxes optimised using nonlinear analysis will be lighter due to the reduction in stresses caused by nonlinear span shortening. This prediction is less certain as it remains unclear whether geometric nonlinearity will allow the optimiser to achieve more or less passive aeroelastic tailoring.

Ribs will show the greatest difference in sizing between linear and nonlinear optimised designs due to the presence of nonlinear Brazier loads.

5

Optimisation Studies

In this chapter I finally apply the developed geometrically nonlinear structural and aerostructural capabilities to structural and aeroelastic design optimisations. In doing so I aim to assess, first, the increase in computational resources required to perform design optimisation with nonlinear structural analysis and, secondly, how geometrically nonlinear phenomena affect the optimal structural and aeroelastic design of high-aspect-ratio wing (HARW). In all optimisations I continue to use the uCRMs, for reasons already described.

Section 5.1 covers a series structural optimisations performed under the same fixed aerodynamic loads used in section 4.2. I perform a baseline optimisation and present the results before highlighting some potential issues in the results, performing further optimisation studies to investigate the issues further. Next I perform aeroelastic optimisation of the uCRMs, using an identical problem formulation but utilising coupled aerostructural analysis and adjoint gradients which allows the optimiser the opportunity to aeroelastically tailor the wing structure to favourably influence the distribution of aerodynamic loads. Finally, in section 5.3 I summarise the findings of the optimisation studies.

5.1. Structural Optimisation Studies

5.1.1. Baseline structural optimisation

The first optimisation study I perform aims to mimic the formulation used by Brooks et al. [112] in their latest publication on the aerostructural optimisation of the uCRM-9. It should be noted that the results of these optimisations are not directly comparable to Brooks et al.'s work due to the slight difference in stiffener orientations I use and the fact that the aerodynamic loading used in these cases is generated on the fine CFD mesh whereas Brooks et al. use the coarse mesh for their work.

In all optimisation cases presented herein, the objective function to be minimised is the mass of the wingbox.

Design Parameterisation

In total, the uCRM-9 and 13.5 wingboxes are parameterised using 586 and 694 design variables respectively. Figure 5.1 and table 5.1 summarise these parameterisations. Each panel in the wingbox is parameterised by four design variables controlling the thickness of the panel skin and the thickness, height and pitch of the smeared panel stiffeners as described in section 2.2.2. In this parameterisation, I give each panel in the upper and lower skins its ow local design variables for the skin thickness, stiffener thickness and stiffener height, while the stiffener pitch across all panels on each skin is assigned to a single global design variable. Similarly, I give each rib and spar panel local skin and stiffener thickness variables, while stiffener height and pitch are controlled globally.

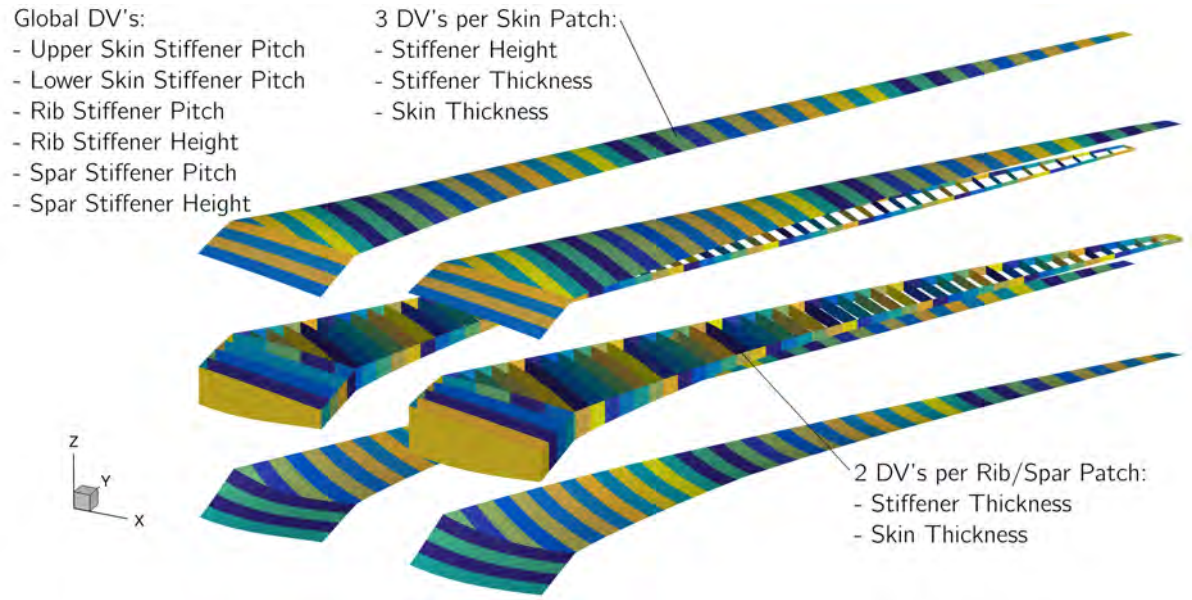

Figure 5.1: Baseline structural parameterisation

Table 5.1: Summary of the design variables used for the baseline uCRM structural optimisation studies

		Components				Quantity		Bounds		
Variable	**Description**	Lower Skin	Upper Skin	Ribs	Spars	uCRM-9	uCRM-13.5	Lower	Upper	**Scaling Factor**
t_{skin}	Local Skin Thicknesses	✓	✓	✓	✓	242	287	3 mm	250 mm	10000
t_{stiff}	Local Stiffener Thicknesses	✓	✓	✓	✓	242	287	2.5 mm	250 mm	10000
h_{stiff}	Local Stiffener Heights	✓	✓			96	114	25 mm	200 mm	1000
λ_{stiff}	Global Stiffener Pitch	✓	✓	✓	✓	4	4	150 mm	300 mm	10
h_{stiff}	Global Stiffener Height			✓	✓	2	2	25 mm	200 mm	1000
					Total	**586**	**694**			

Constraints

I apply constraints on the Von Mises buckling failure criteria in the 2.5 g, -1 g and 1 g 'gust' manoeuvre conditions described in section 4.1.2. I split the wingbox into four separate regions, upper skin, lower skin, ribs and engine mount, and spars, within which, the element Von Mises and buckling values are combined into a single constraint values using KS aggregation. In the 2.5 and 1 g load cases I apply Von Mises constraints to all regions and buckling constraints on all regions except the lower skin. In the -1 g case I constrain only the lower skin buckling value. As explained in section 4.1.2, I apply a safety factor of 1.5 to the failure constraints in the 2.5 and -1 g loadcases a safety factor of 2.67 to the 1 g constraints. This gives a total of 15 failure and buckling constraints.

To enforce a realistic structural design, I apply a series of constraints to the panel design variables. On all panels, I constrain the stiffener and skin thicknesses to be within 2.5 mm of each other. Additionally, to avoid abrupt changes in panel sizing I apply adjacency constraints to limit the change in skin thickness, stiffener thickness and stiffener height between adjacent skin panels to be $\leq$0.5 mm and $\leq$2 mm on adjacent spar panels. These constraints, although numerous, are linear and are therefore handled efficiently by the optimiser used for these studies, SNOPT, without the need for repeated gradient calculations.

Optimiser

I use SNOPT [113], a high performance sequential quadratic programming (SQP) optimiser well suited to large, sparse, constrained optimisation problems. SNOPT handles linear and nonlinear constraints

Table 5.2: Constraints applied to the baseline structural optimisation problem.

		Components				Quantity	
Constraint	**Description**	Lower Skin	Upper Skin	Ribs	Spars	uCRM-9	uCRM-13.5
$1.5KS_{vm,2.5g}$	2.5 g Von Mises failure	✓	✓	✓	✓	4	4
$1.5KS_{buckling,2.5g}$	2.5 g buckling failure		✓	✓	✓	3	3
$2.67KS_{vm,1g}$	1 g Von Mises failure	✓	✓	✓	✓	4	4
$2.67KS_{buckling,1g}$	1 g buckling failure		✓	✓	✓	3	3
$1.5KS_{buckling,-1g}$	-1 g buckling failure	✓				1	1
$t_{skin,i} - t_{stiff,i}$	Stiffener-skin thickness difference	✓	✓	✓	✓	242	287
$t_{skin,i} - t_{skin,i+1}$	Skin thickness adjacency	✓	✓		✓	183	219
$t_{stiff,i} - t_{stiff,i+1}$	Stiffener thickness adjacency	✓	✓		✓	183	219
$h_{stiff,i} - h_{stiff,i+1}$	Stiffener height adjacency	✓	✓			92	110
					Total	**715**	**850**

separately and, after satisfying them in the initial iteration, guarantees feasibility of all linear constraints at every subsequent iteration. In all optimisations I use feasibility and optimality tolerances of 10^{-7} these correspond to a maximum constraint violation of around 10^{-5} and a reduction in the norm of the Lagrangian of 4-5 orders of magnitude.

Results

Table 5.3 shows the number of optimiser iterations and the runtime for each of the baseline structural optimisations on 16 cores of a 3.0 GHz Intel Xeon Gold 6154 processor. Performing structural optimisation with nonlinear analysis appears to increase the computational cost of each optimiser by approximately a factor of 10. In the case of the uCRM-13.5 this increase is slightly reduced as the nonlinear optimisation converges in 24% fewer iterations, the runtime is however still almost an order of magnitude greater than the linear case.

Table 5.3: Comparison of time and number of optimiser iterations required for each baseline structural optimisation.

	uCRM-13.5			**uCRM-9**			
Value	Linear	Nonlinear	Δ(%)	Linear	Nonlinear	Δ(%)	**Units**
Major Iterations	441	336	-24	485	475	-2	
Function Evaluations	1084	841	-22	1087	1044	-4	
Time	1497	13060	772	1728	22702	1214	s
Time per Major iteration	3.4	38.9	1045	3.6	47.8	1242	s
Time per Function Evaluation	1.4	15.5	1024	1.6	21.7	1268	s

The following graphs compare the results of the linear and nonlinear structural optimisations of the uCRM-13.5. Figure 5.2 compares the histories of the feasibility, optimality and total wingbox mass for each optimisation and figure 5.3 compares the histories of the masses of each component group.

Both the linear and nonlinear optimisations achieve a significant decrease in mass over the baseline uCRM wingbox design and satisfy both the optimality and feasibility convergence criteria. The wingbox optimised with nonlinear analysis is 6.1% heavier than the linear optimised design, an increase that is consistent with

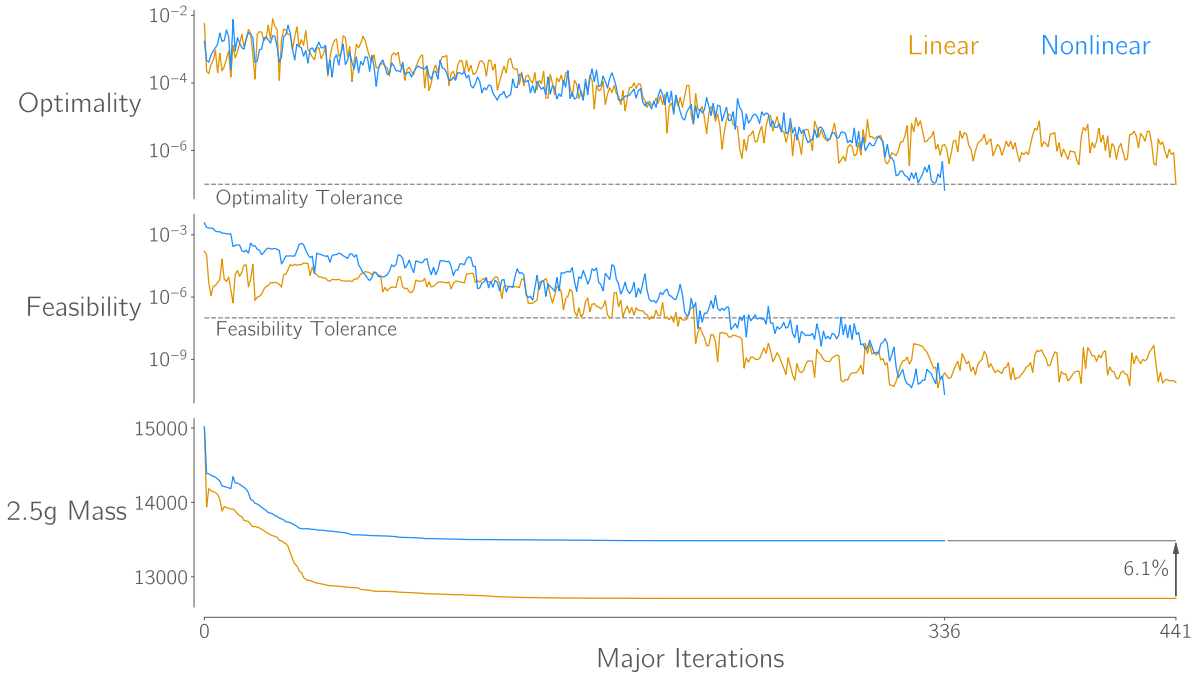

Figure 5.2: Structural optimisation of the uCRM-13.5 with nonlinear structural analysis converges in fewer iterations but results in a 6.5% greater mass than with linear analysis.

the increased bending stresses observed in the structural analysis studies. This mass increase from linear to nonlinear optimised wingbox is greatest in the lower skin, followed by the upper skin, ribs and then spars, again suggesting that the cause of the mass increase is primarily due to bending stresses.

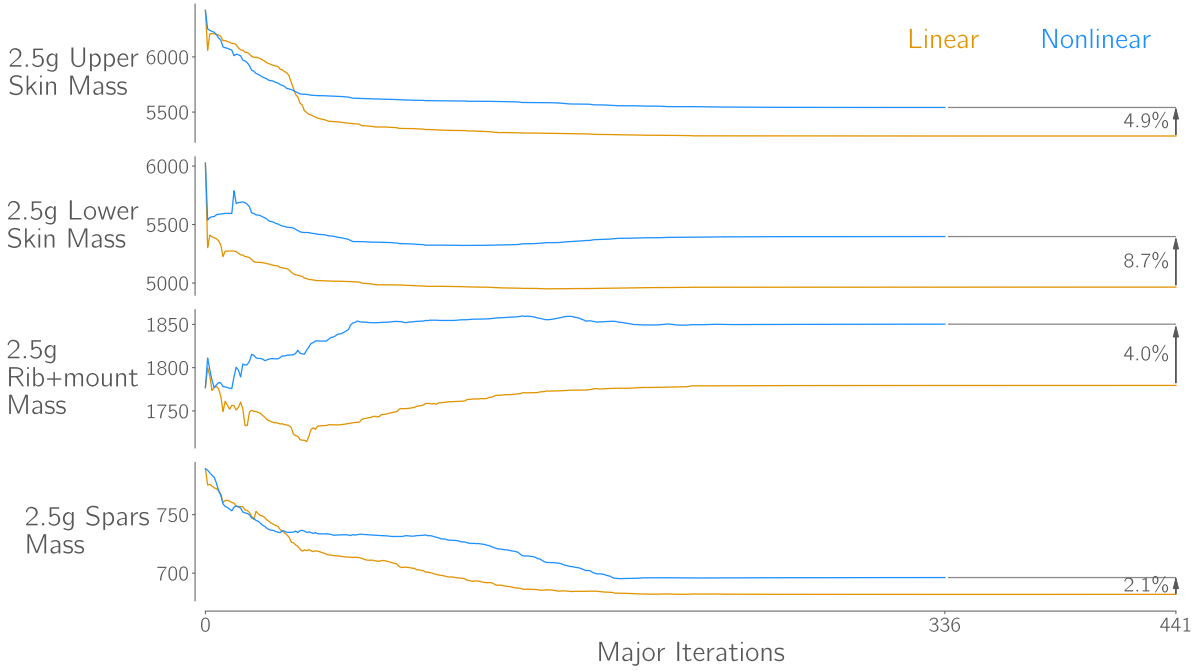

Figure 5.3: The majority of the mass difference between linear and nonlinear optimised designs is in the upper and lower skins.

In figures 5.4 I plot the spanwise distribution of the structural sizing of each wingbox component group. In order to effectively display the effective stiffness of each component, I plot what I refer to as aggregated sizing values, which combine the stiffness contribution of both the panel skins and stiffeners into a single

value. The 'axial effective thickness' is the thickness of an equivalent unstiffened panel with the same axial stiffness (in the stiffener-wise direction) as the stiffened panel and the 'bending effective stiffness' follows the same rationale but for bending stiffness. Note that the axial effective thickness is directly proportional to the cross-sectional area of each panel and can thus also be seen as a proxy for the mass of each panel.

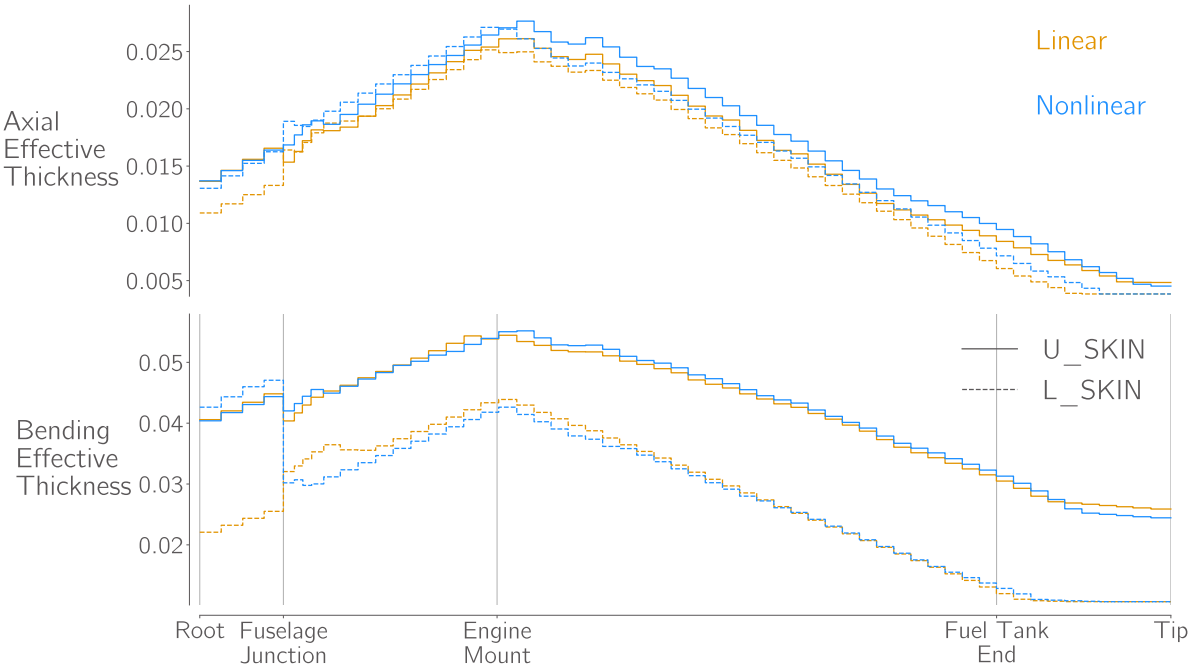

Figure 5.4: The increase in skin stiffness in nonlinear optimised design is almost uniform, similar to increase in stress seen in the structural analysis studies.

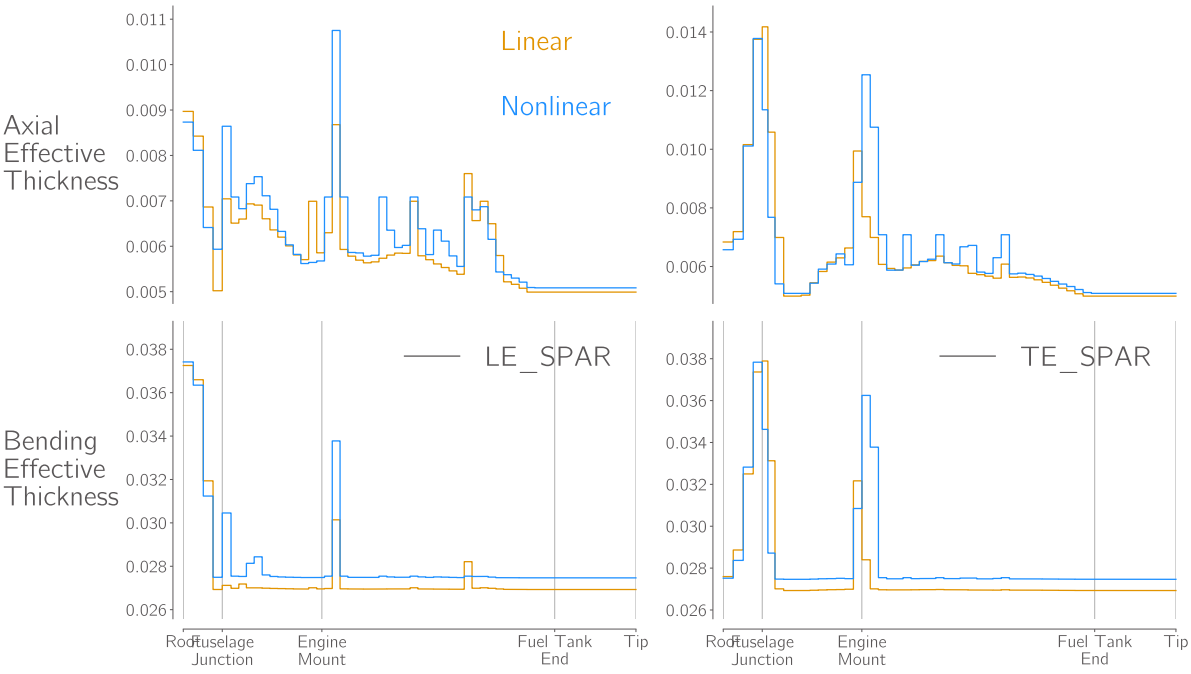

Figure 5.5: Spar sizing is highly effected by external mass loads in the nonlinear case, bending stiffness also increased to resist buckling.

The increase in the axial effective thickness of the upper and lower skins from linear to nonlinear optimised designs is relatively uniform along the entire wingspan which is again consistent with the uniform increase

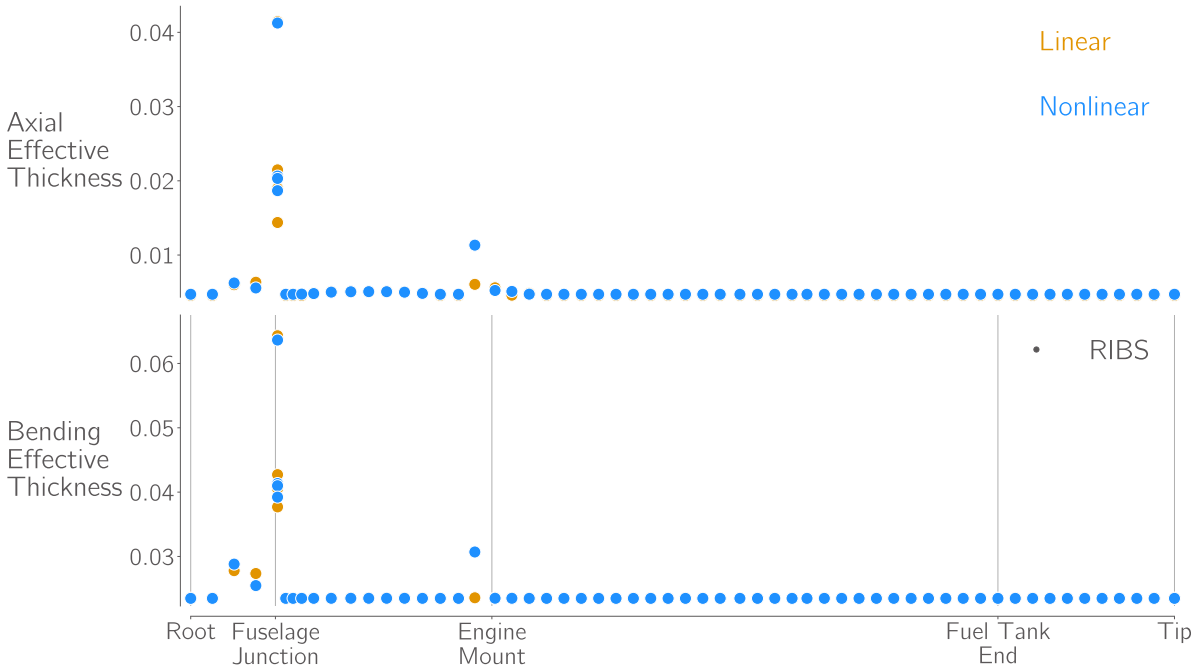

Figure 5.6: Brazier loading has no clear effect on rib sizing due to limited parameterisation.

in bending stresses observed in the structural analysis studies. The spanwise variation of axial effective thickness is almost perfectly linear, rising from the wing root to engine mount and falling from engine mount to wingtip. This suggests that the rate of reduction reduction of the thickness of the skins away from the engine mount region is limited by the applied adjacency constraints and that further mass savings could be made by loosening them. There is significantly less difference in the bending stiffness distributions of the linear and nonlinear optimised designs, with the bending stiffness of the upper skins being almost equal over large portions of the span. On the lower skin, more significant differences are observed, the bending stiffness of the linear optimised lower skin is significantly higher over the portion of the wing between the fuselage junction and engine mounting location, perhaps indicating that the compressive stresses seen in the -1 g manoeuvre condition are greater in the linear case. In contrast, the bending stiffness of the lower skin of the centre wingbox (inside the fuselage) is more than $8\times$[1] greater in the nonlinear optimised design.

Both the leading and trailing edge spars are thicker in the nonlinear case over the majority of the wing span. There are a series of prominent spikes in the effective axial thickness of both spars in the nonlinear case which correspond to the spar panels subject to loads from the external leading and trailing edge masses. These reinforcements are then a symptom of the previously discussed issues with the application of these external mass loads which causes unrealistic out of plane loading on the spar panels in the nonlinear analyses. Also noticeable is an almost completely uniform increase in bending stiffness in both the nonlinear optimised spars, indicating a uniform increase in buckling loads. This is likely due to an increase in the global spar stiffener height variable. Without further analysis of the critical buckling loads in the spars it is not possible to draw any conclusions on whether this increase is driven by a uniform increase in buckling loads over the spar or simply by the sizing of one of the more highly loaded spar panels around the wing root or engine mount.

Despite earlier evidence of significant Brazier loading increasing the rib buckling loads in the nonlinear analyses, the majority of the ribs are sized identically in the linear and nonlinear analyses. The likely explanation for this is that the stiffener pitch and height variables, shared by all ribs in this parameterisation are driven by the fuselage junction rib, which is loaded drastically more highly than any of the other ribs. As a result, the stiffeners on the remaining ribs are oversized for the loads they experience in both the linear

[1]The bending effective thickness is only greater by a factor of 2 but bending stiffness is proportional to the cube of thickness.

and nonlinear analyses. The only exception to this rule is the rib adjacent to the engine mount which is significantly stiffer in the nonlinear optimised design. This again is likely due to the geometric inaccuracies in the application of the engine loads in the nonlinear analysis.

Below, I present the same graphs as above for the linear and nonlinear structural optimisations of the uCRM-9.

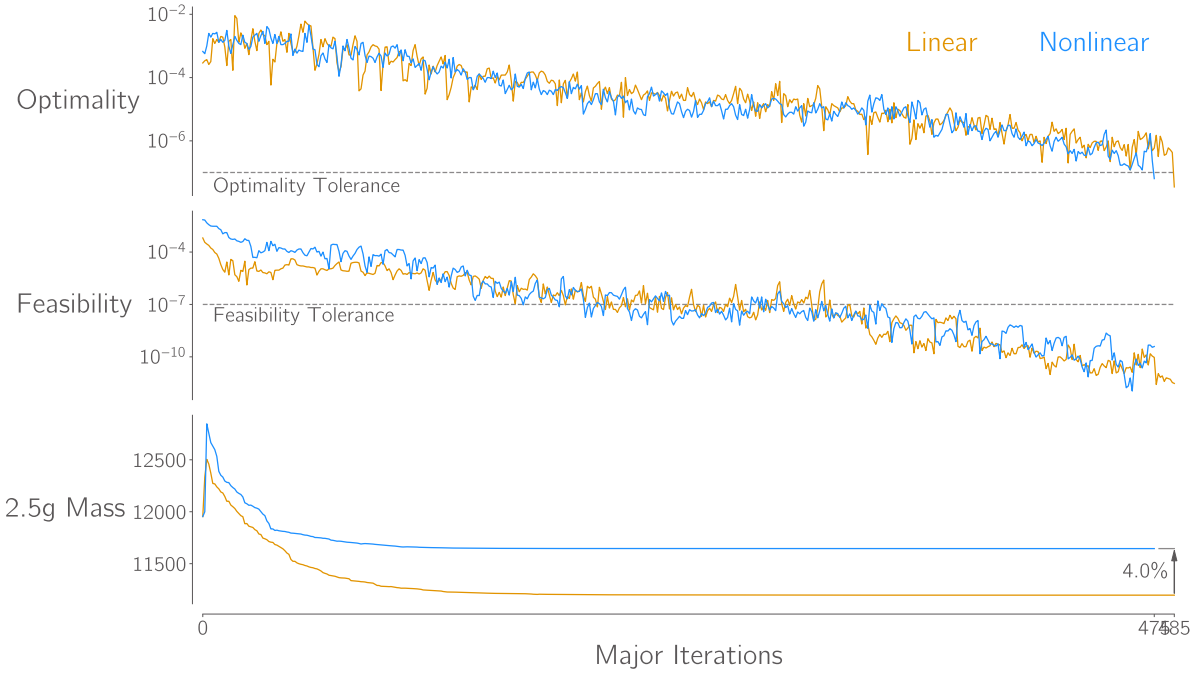

Figure 5.7: The mass increase due to structural nonlinearity is the uCRM-9 is smaller than with uCRM-13.5 but still significant.

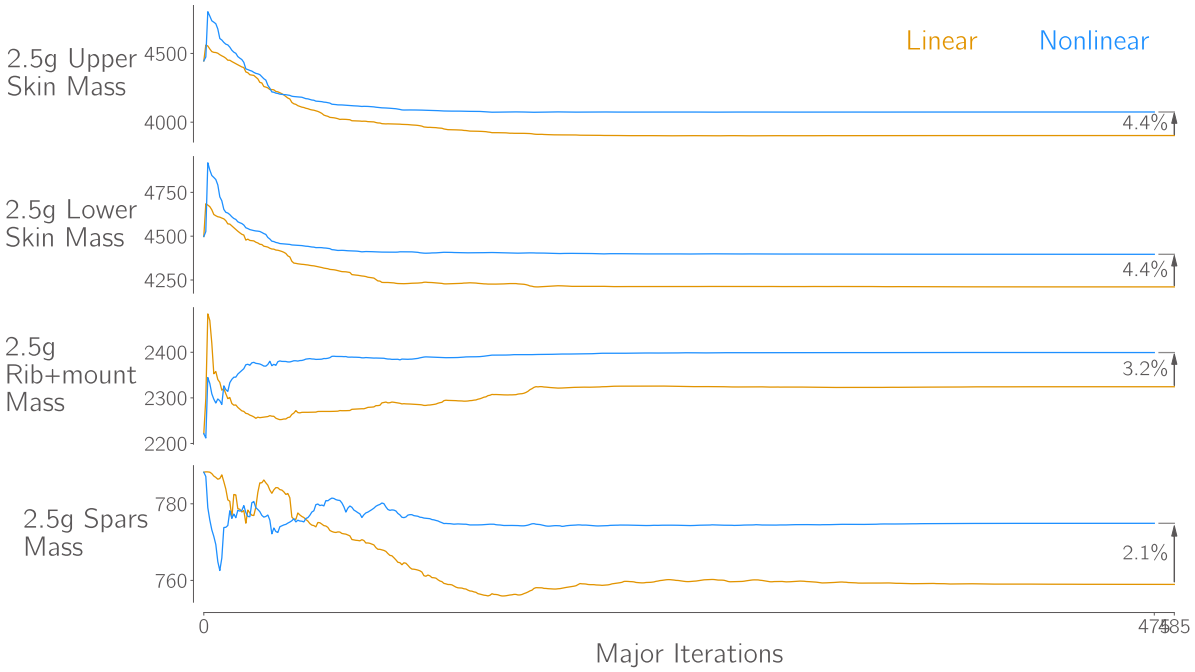

Figure 5.8: Again, majority of the mass difference between linear and nonlinear optimised uCRM-9 designs is in the upper and lower skins.

In the uCRM-9, the differences between linear and nonlinear and nonlinear optimised designs are similar

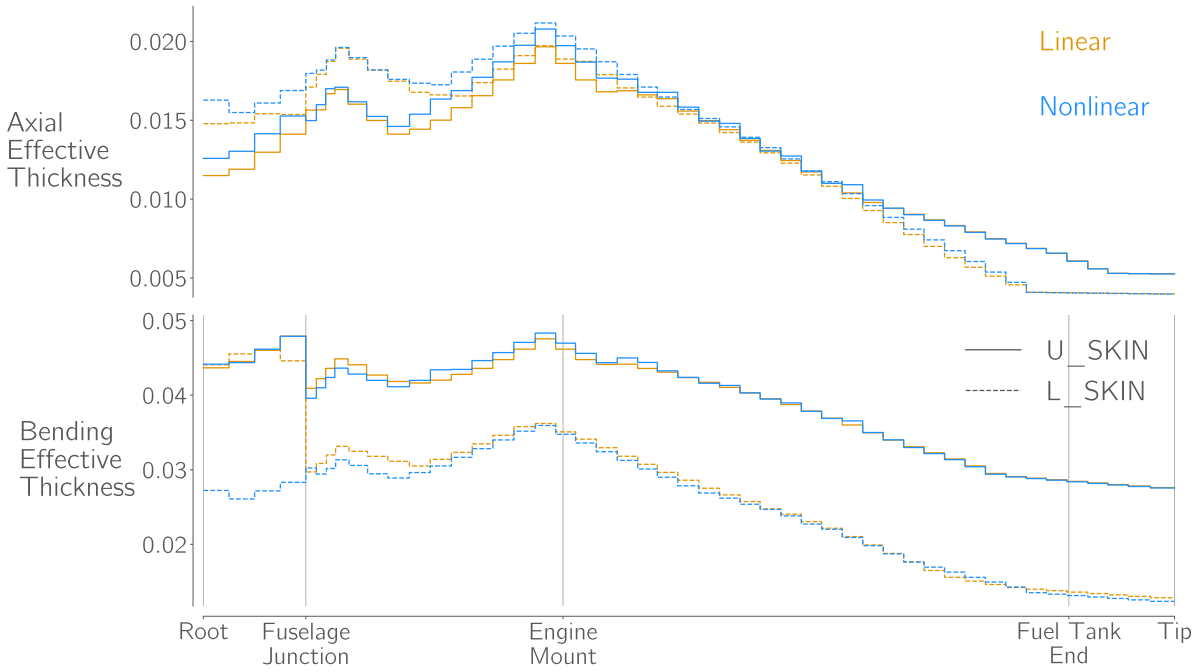

Figure 5.9: The increase in skin stiffness in nonlinear optimised design is less uniform in the uCRM-9 than the uCRM-13.5. Most of the additional stiffness is added near wing root.

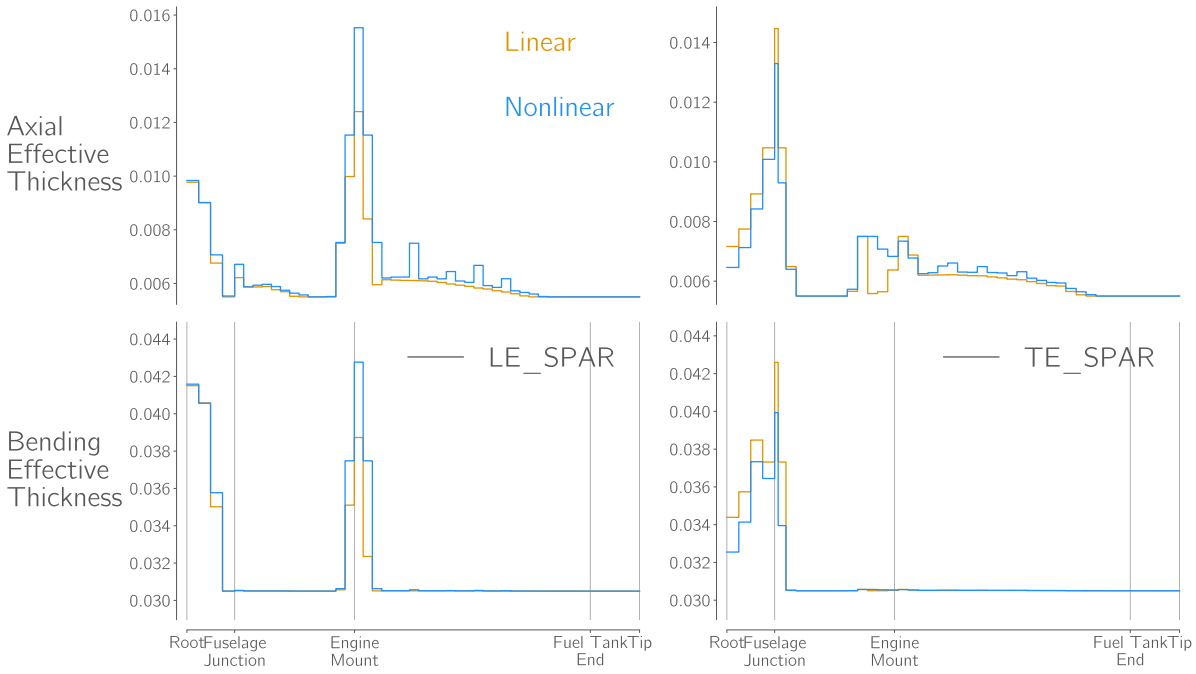

Figure 5.10: Spar sizing in the uCRM-9 is again highly effected by external mass loads in the nonlinear case but no significant increase in bending stiffness indicating spars are less buckling critical in uCRM-9.

and slightly smaller than those seen in the uCRM-13.5, but still significant. The mass difference between the linear and nonlinear optimised designs reduces to 4% and again, the upper and lower skins see the largest changes in mass, followed by the ribs and then spars.

The increase in the thickness of the upper and lower skins in the nonlinear optimised design is less uniformly distributed in the uCRM-9 than in the uCRM-13.5 with the greatest differences seen in the centre wingbox and the areas surrounding the engine mount. The linear and nonlinear optimised spar sizing distributions

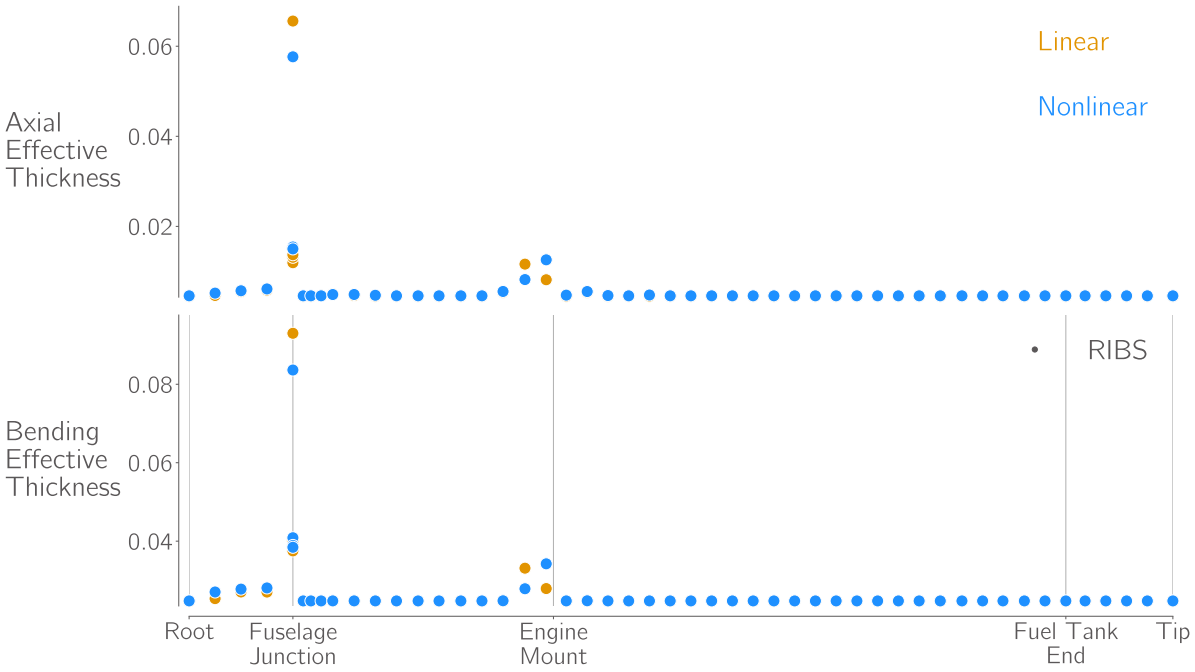

Figure 5.11: Again, Brazier loading has no clear effect on rib sizing in the uCRM-9 due to limited parameterisation.

are more similar than in the uCRM-13.5. The thickness peaks due to external mass loads are only present in the leading edge spar and the uniform increase in bending stiffness seen in the nonlinear optimised uCRM-13.5 design is not seen. Once again, the majority of the ribs are oversized due to the parameterisation issues described above and there is consequently no difference in the sizing of most of the ribs.

Interestingly, same large discrepancy in the bending stiffness of the lower skin of the centre wingbox is seen in the uCRM-9 designs, but in this case the linear optimised design is the one being significantly reinforced. To explore this phenomena, figure 5.12 shows the buckling criteria distributions in all 4 optimised designs. However, viewing these results, the rationale for the changes remains somewhat unclear as in both designs where the centre wingbox lower skin is strongly reinforced, it is reinforced far beyond the level required to resist buckling. It is possible therefore that the optimiser chooses this drastically increased bending stiffness in order to reduce the stresses in surrounding regions of the wingbox, rather than reducing the stress in the centre wingbox lower skin itself.

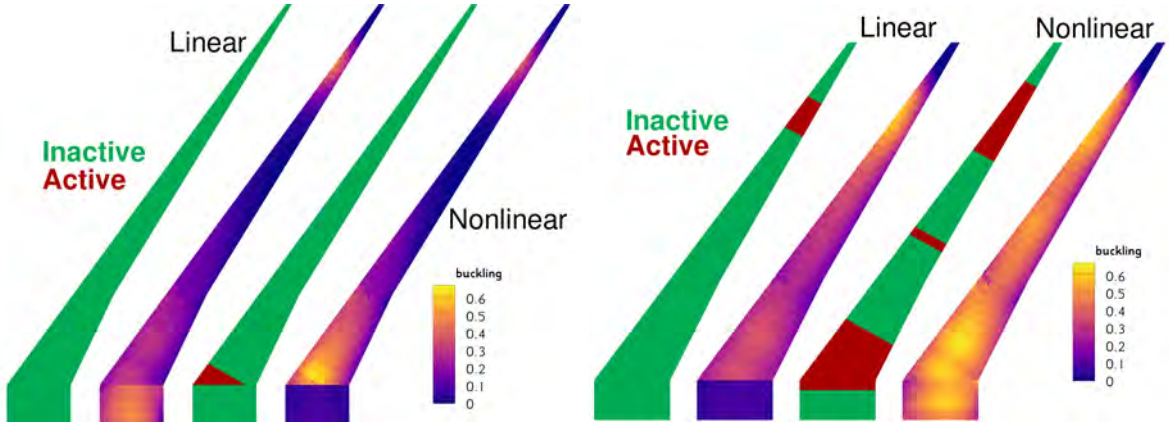

(a) uCRM-13.5.

(b) uCRM-9.

Figure 5.12: Comparison of lower skin buckling constraint values from the -1 g loadcase in the linear and nonlinear optimised uCRM wingboxes.

5.1.2. Optimisation with Greater Design Freedom

I now run a set of optimisations with increased design freedom, allowing the optimiser to correctly size each of the rib and spar panels in order to assess the effect of Brazier loading on wing mass. Figure 5.13 summarises the new design parameterisation, in which each rib and spar panel is given it's own stiffener height and pitch design variable. Additionally, I loosen the adjacency constraints to 2.5 mm, this value is based on values used in previous works [61, 110] and these loosened adjacency constraints are also applied to the new local rib and spar stiffener design variables. This change is motivated by the fact that the majority of the adjacency constraints in the previous optimisation were active, heavily restricting the optimiser and resulting in the primarily linear spanwise variations in the sizing of the upper and lower skins.

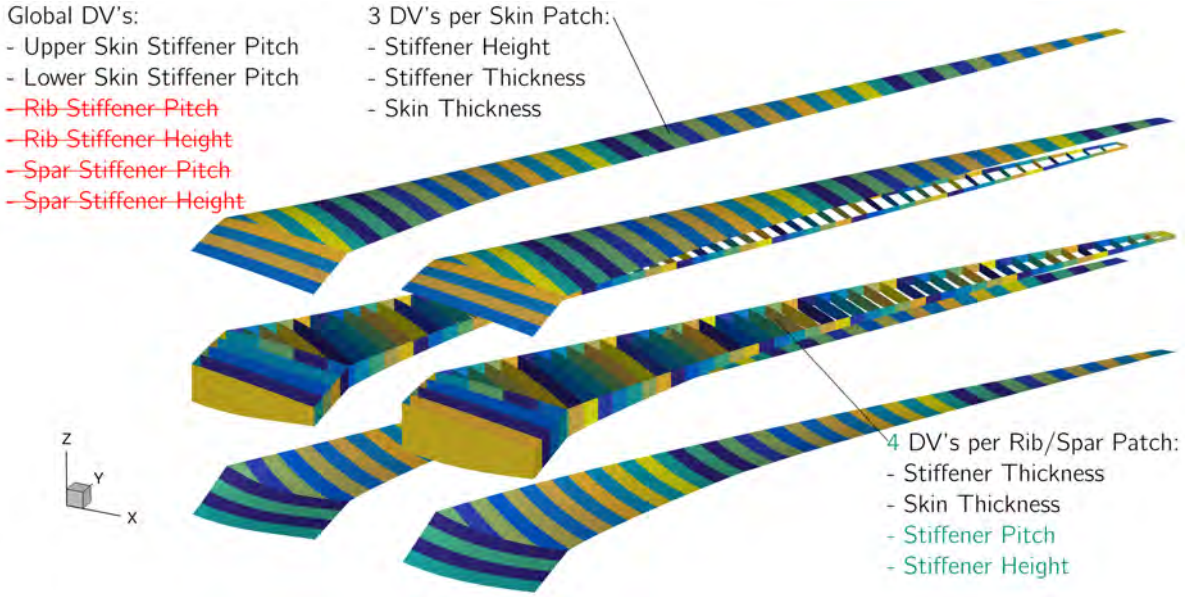

Figure 5.13: Structural parameterisation with separate stiffener design variables for each rib and spar patch.

Results

All results presented below are for the uCRM-13.5, results from the uCRM-9 optimisations are not shown as they follow the same trends. Figures 5.14 and 5.15 compare the optimisation histories of the baseline and increased freedom optimisations for the linear and nonlinear structural formulations respectively, while figures 5.16 and 5.17 compare the histories of the linear and nonlinear increased design freedom optimisations.

All of the increased design freedom optimisation reach the specified feasibility tolerance but none reach the required optimality value before reaching the set major iterations limit. However, looking at the mass histories of each case it is clear that little change in the designs would be seen had the optimisations continued. The additional design freedom leads to a 7.9% decrease in mass in the nonlinear case and 5.6% decrease in the linear case. As a result, the difference between the linear and nonlinear optimised designs decreases to 3.8%. As expected, the additional design freedom given to the ribs and spars means they now show larger relative mass differences between linear and nonlinear optimised designs, with the ribs becoming the component group that is effected most by the nonlinearity.

Figure 5.18 compares the rib sizing distributions from the nonlinear baseline and increased design freedom optimisations. The oversizing of the ribs in the baseline case is clearly demonstrated by the fact that the bending stiffness of most of the ribs is reduced by an order of magnitude in the increased design freedom case. In figure 5.19, the impact of Brazier load on the rib sizing are clear, with the ribs between the fuselage junction and engine mount having noticeably greater bending stiffness in the nonlinear optimised design. It is worth noting, however, that the axial effective thickness of these ribs does not increase significantly,

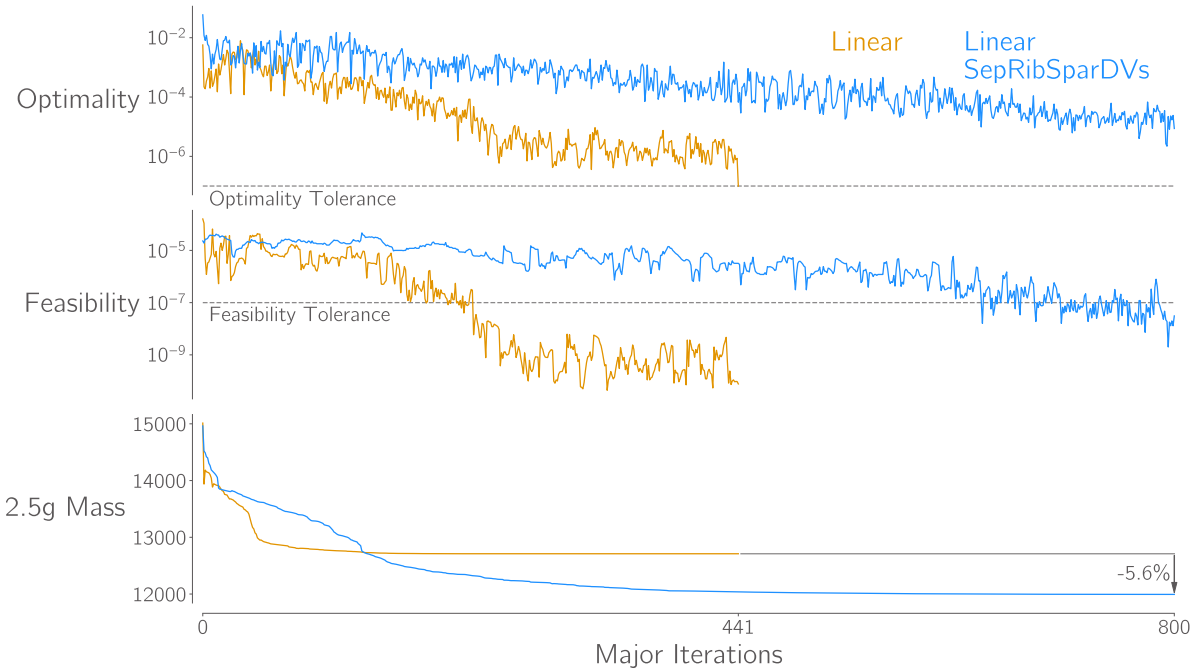

Figure 5.14: Greater design freedom results in a further 5.6% mass reduction in the linear case

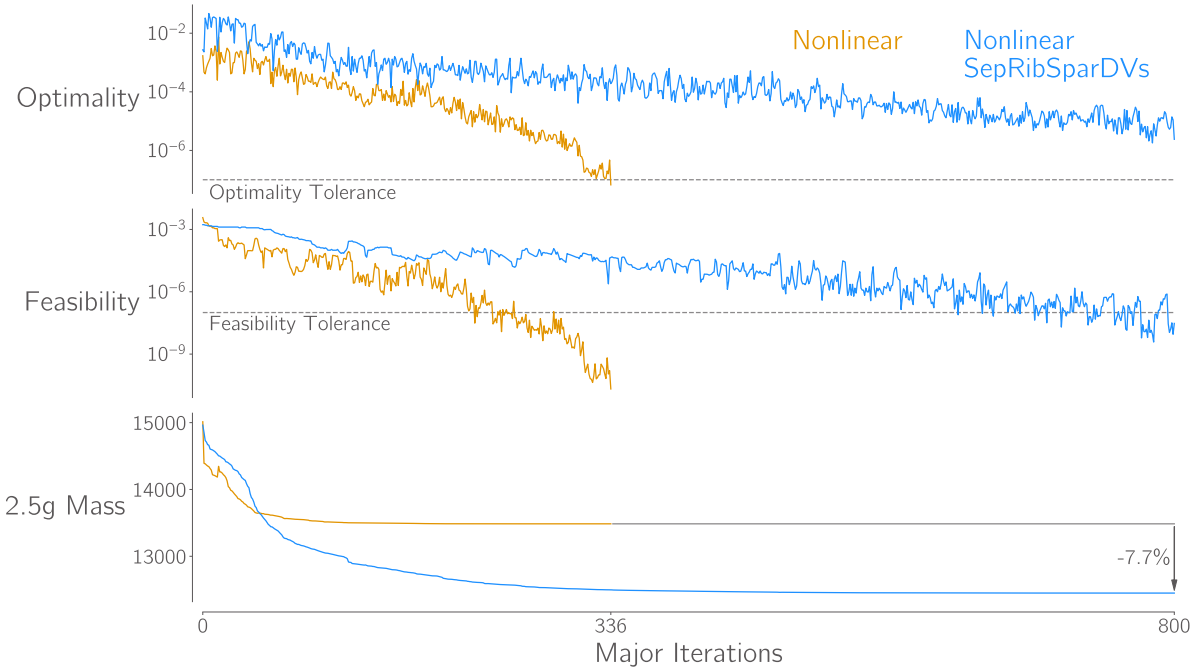

Figure 5.15: Greater design freedom results in a further 7.9% mass reduction in the nonlinear case

meaning that the optimiser chooses to shift material in the rib panels from the skin to the stiffeners, without increasing their overall mass. It is likely that the majority of the difference in rib mass between the linear and nonlinear optimised designs occurs in the fuselage junction rib.

Similar effects are seen in the sizing of the spars. As shown in figure 5.20, giving each spar panel control over its own stiffener design variables allows the optimiser to drastically reduce the spar's stiffness over the majority of the wingspan. This supports the earlier assertion that the increase in spar bending stiffness seen in the nonlinear baseline optimisation was driven by the stiffener sizing required by one of the highly

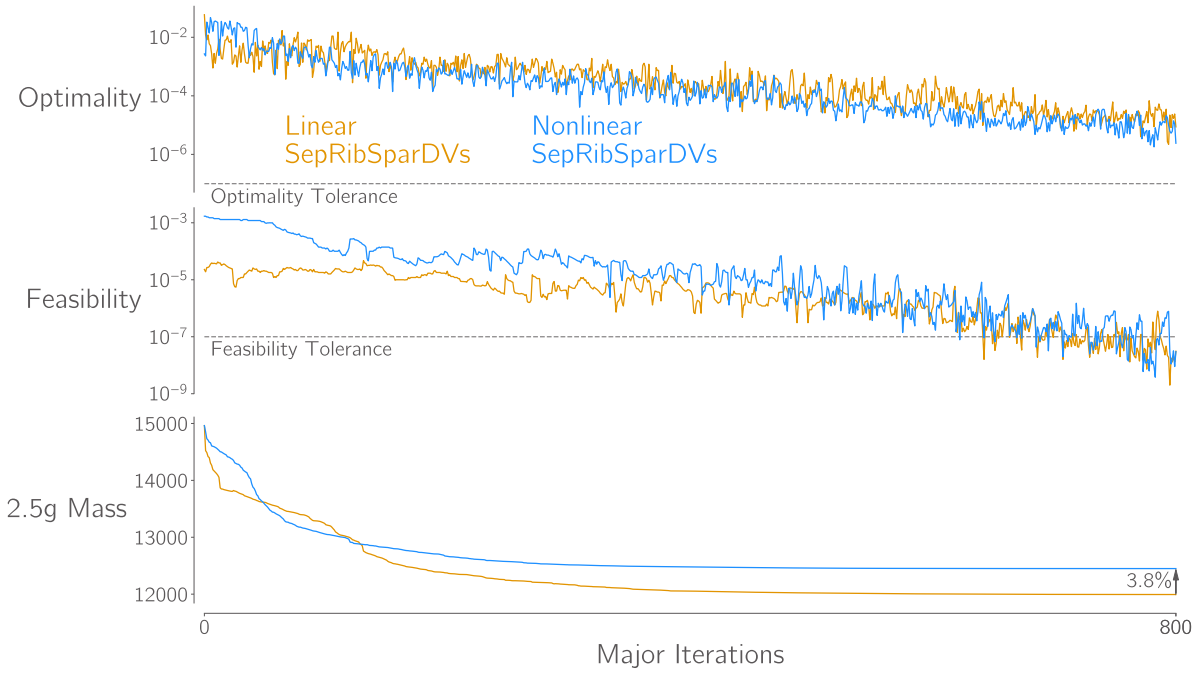

Figure 5.16: Increased design freedom reduces the gap between the masses of the wingboxes optimised with linear and nonlinear analysis

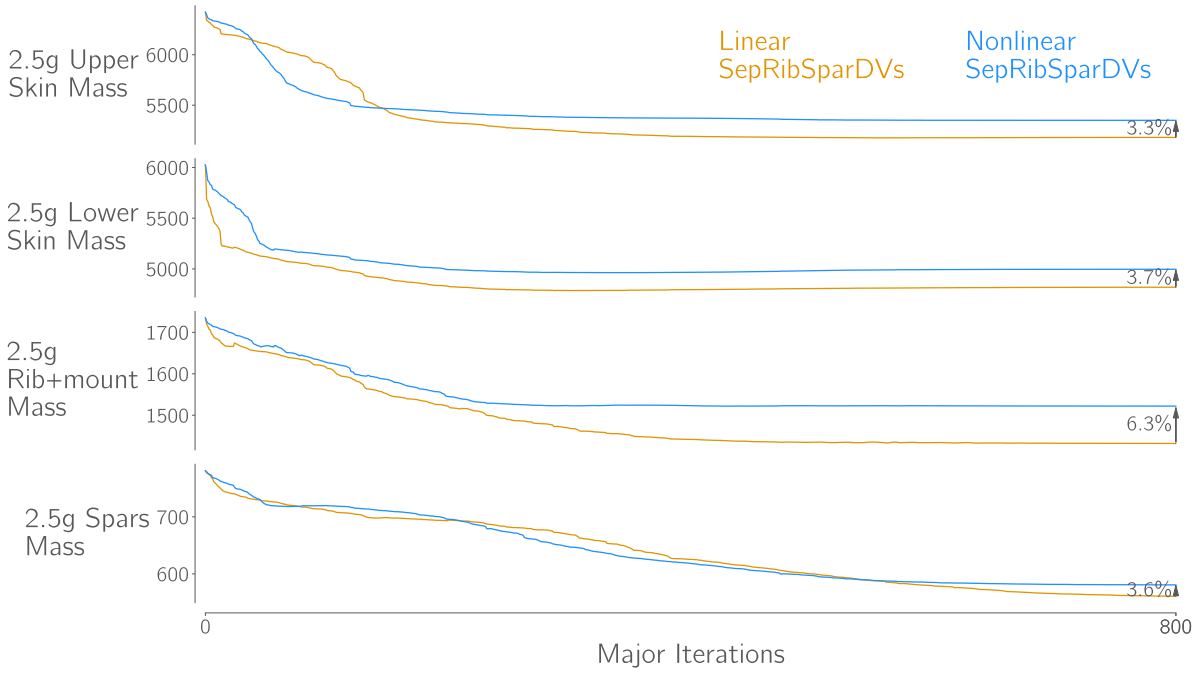

Figure 5.17: With increased design freedom, the ribs and spars now show the largest relative difference in mass between linear and nonlinear optimised designs.

loaded spar panels at the wing root. Figure 5.21 compares the spar sizing from the linear and nonlinear increased design freedom cases, the largest deviations occur in the final 10% of the wingspan, where the linear optimised design requires greater bending stiffness. A similar phenomena is seen in the ribs in this section of the wing but I leave further investigation of the cause to future work.

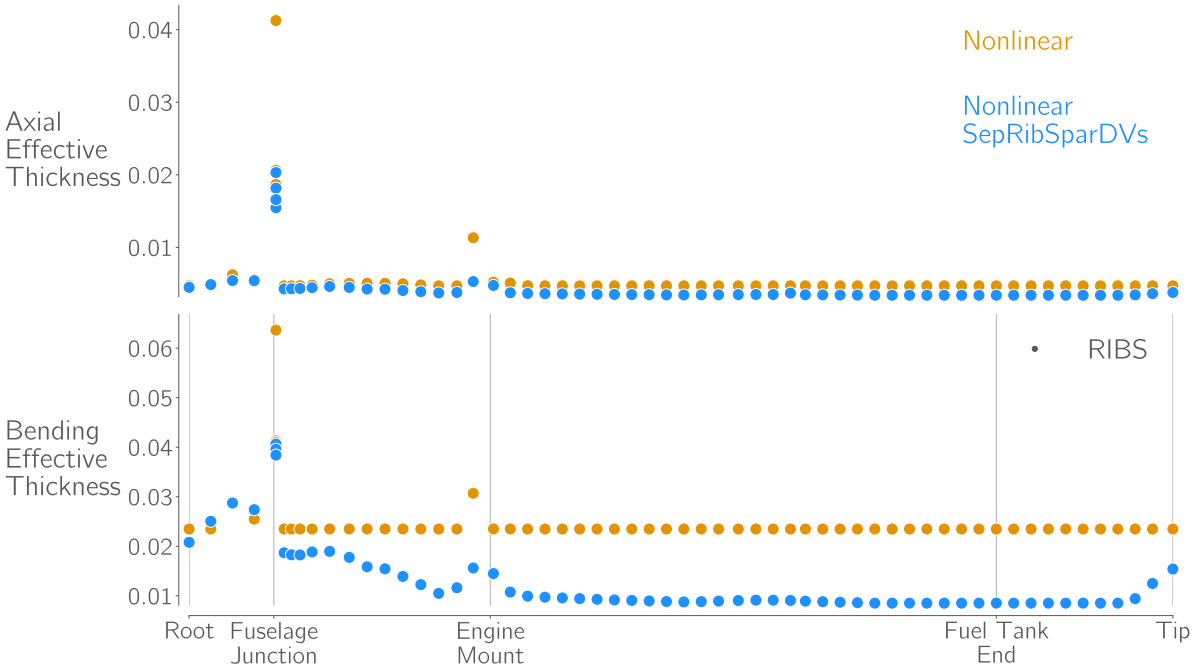

Figure 5.18: When given their own stiffener sizing design variables, the majority of the ribs can be significantly lightened.

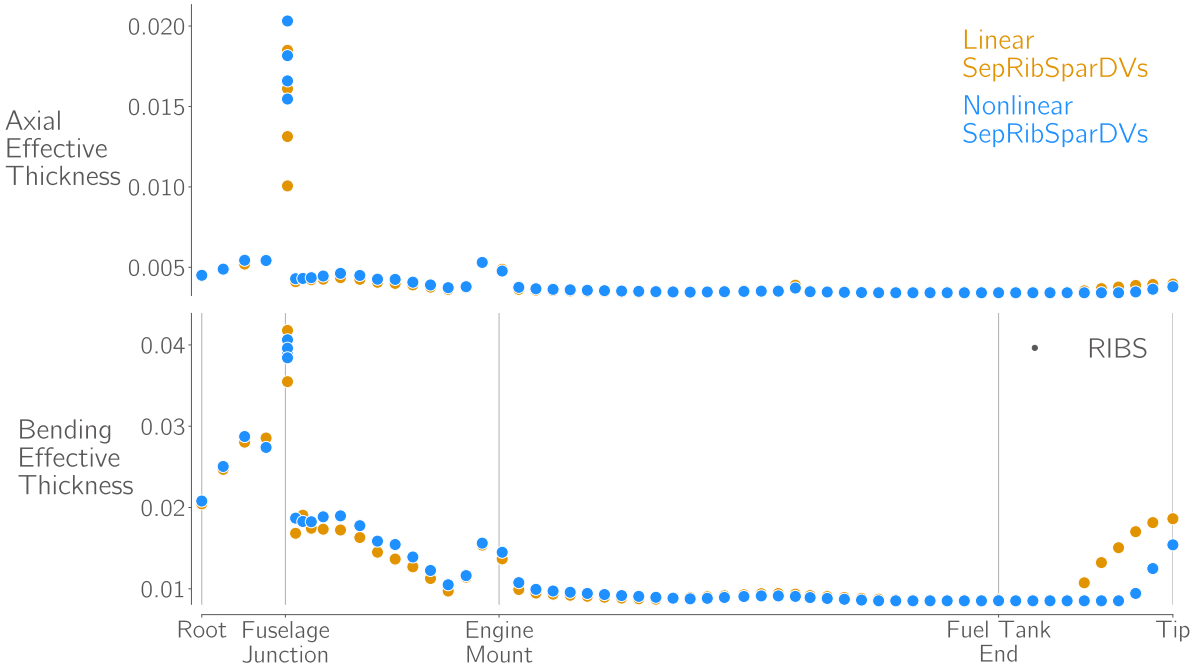

Figure 5.19: Brazier loads require greater bending stiffness in the ribs between the engine mount and fuselage junction.

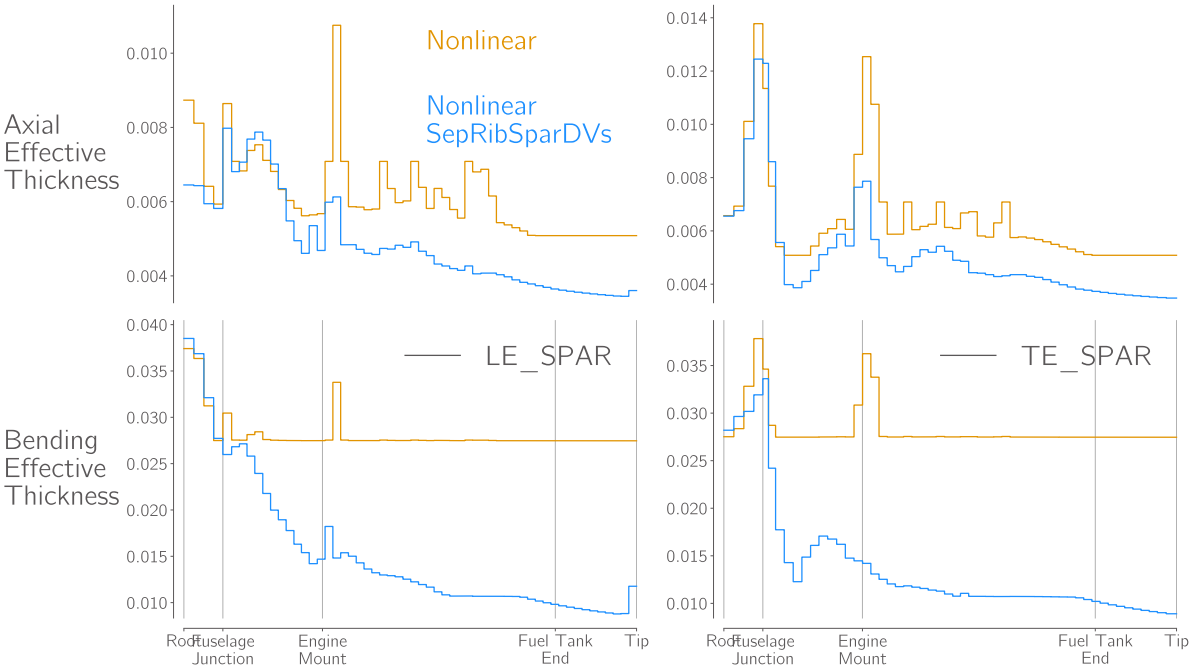

Figure 5.20: Giving each spar panel control over its own stiffener design variables allows the optimiser to drastically reduce the spar's stiffness over the majority of the wingspan.

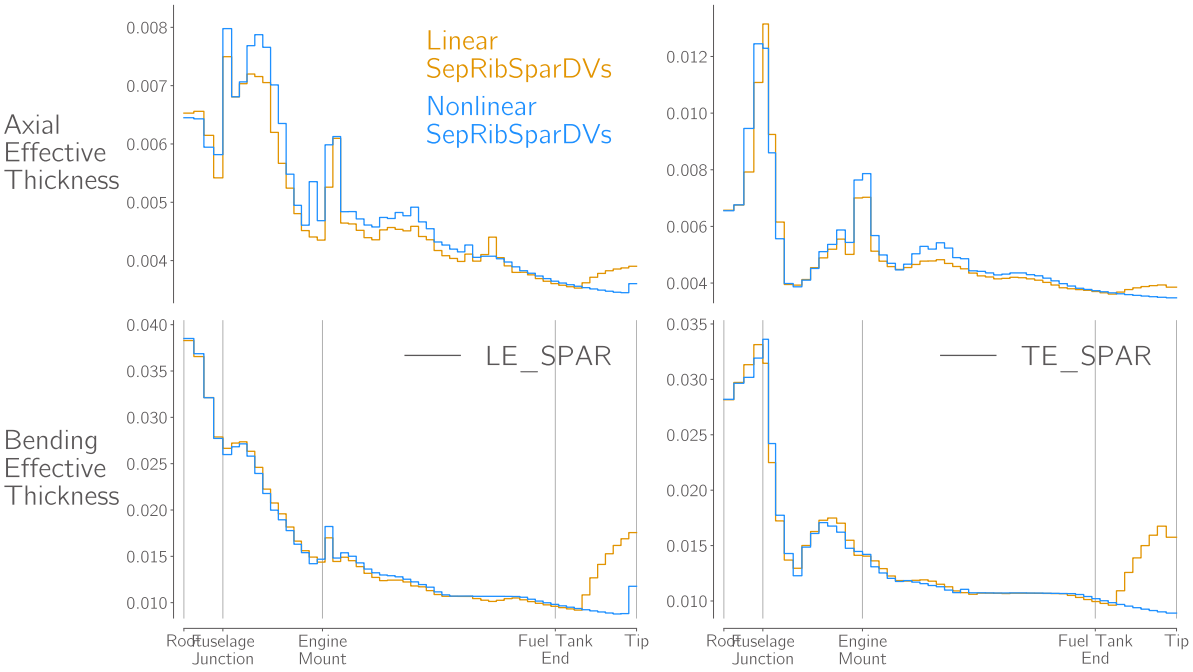

Figure 5.21: As seen in the ribs, the wingbox optimised with linear analysis and increased design freedom requires greater bending stiffness in the spar panels towards the wingtip.

5.1.3. Bi-level Optimisation

I now run two sets of optimisations using a bi-level approach which aims to accelerate the process of achieving a geometrically nonlinear optimised structural design. The approach works by first optimising the wingbox using linear analysis, before using the linear optimised design as the starting point for an optimisation with nonlinear analysis. The motivation of this approach is to use the fast linear analysis to perform the exploratory phase of the optimisation, before the nonlinear analysis is used to fine tune the design for the increased stresses seen in the nonlinear case. I test this approach on the baseline optimisation formulation covered at the beginning of this section, in one case starting from the same reasonably good baseline uCRM wingbox design as the previous optimisations and, in a second case, starting from a poor initial design with uniform sizing throughout the wingbox. This poor initial design is not only heavier than the good starting design used previously, but also violates the failure constraints by a much greater margin, making the optimisation problem significantly more challenging.

Results

Figure 5.22 shows that, when starting from the good initial design, the bi-level approach is 7% slower than optimising with only nonlinear analysis.

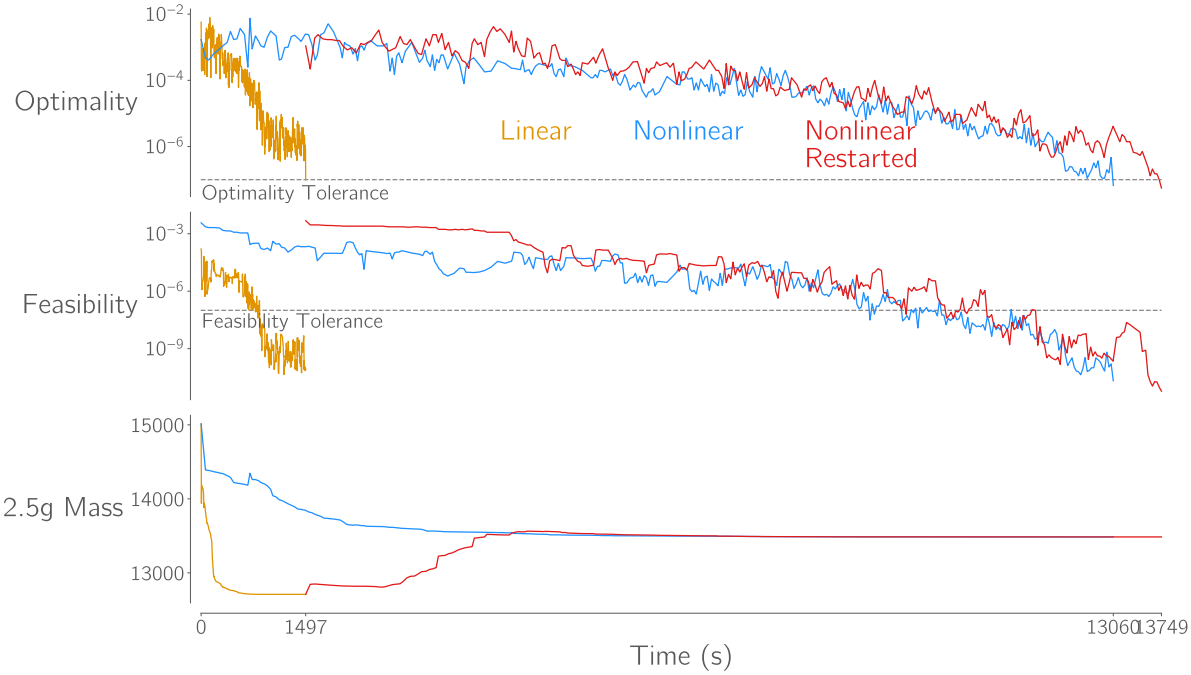

Figure 5.22: When starting from the 'good' baseline uCRM design, the bi-level approach is 7% slower

However, when starting from the poor design, as shown in figure 5.23, the linear phase of the optimisation is able to quickly produce a reasonable design which greatly reduces the number of iterations required in the nonlinear phase. As a result, the bi-level approach is 37% faster.

Crucially, figure 5.24 shows that both the standard and bi-level optimisation approaches converge to the same optimised design, including the large increase in the bending stiffness of the centre wingbox lower skin. The ability of both these optimisation to reach the same optimum from such different initial points suggests that the design space for this optimisation problem is well formed and does not contain local optima.

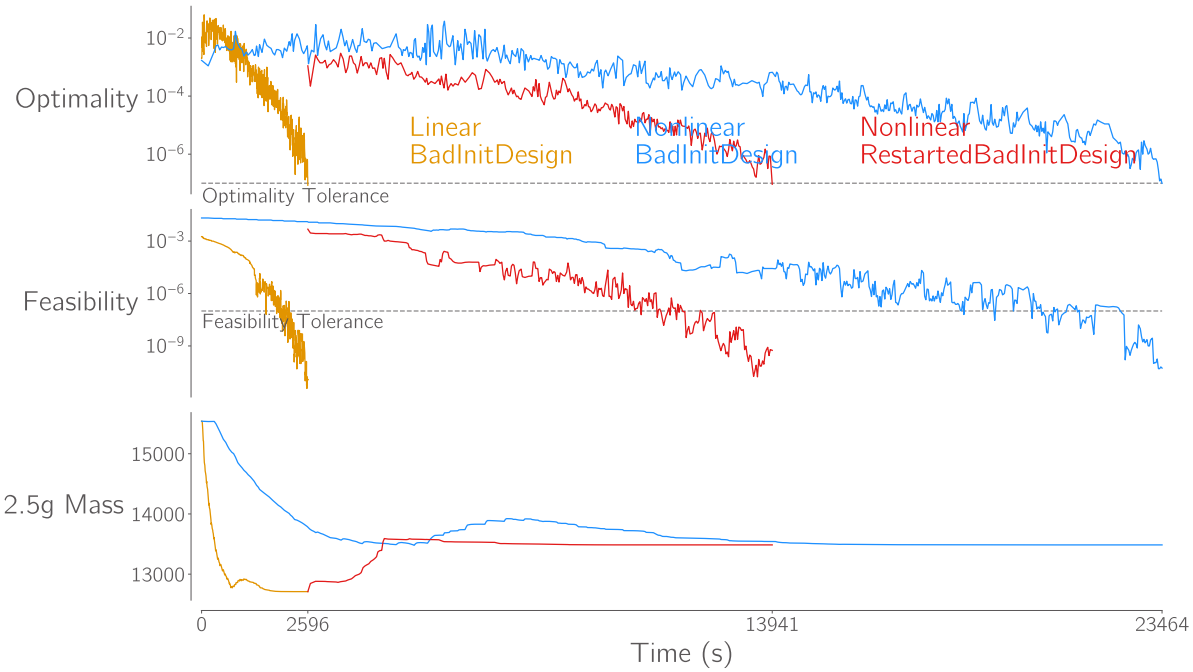

Figure 5.23: When starting from a very poor initial design, the bi-level approach is 37% faster.

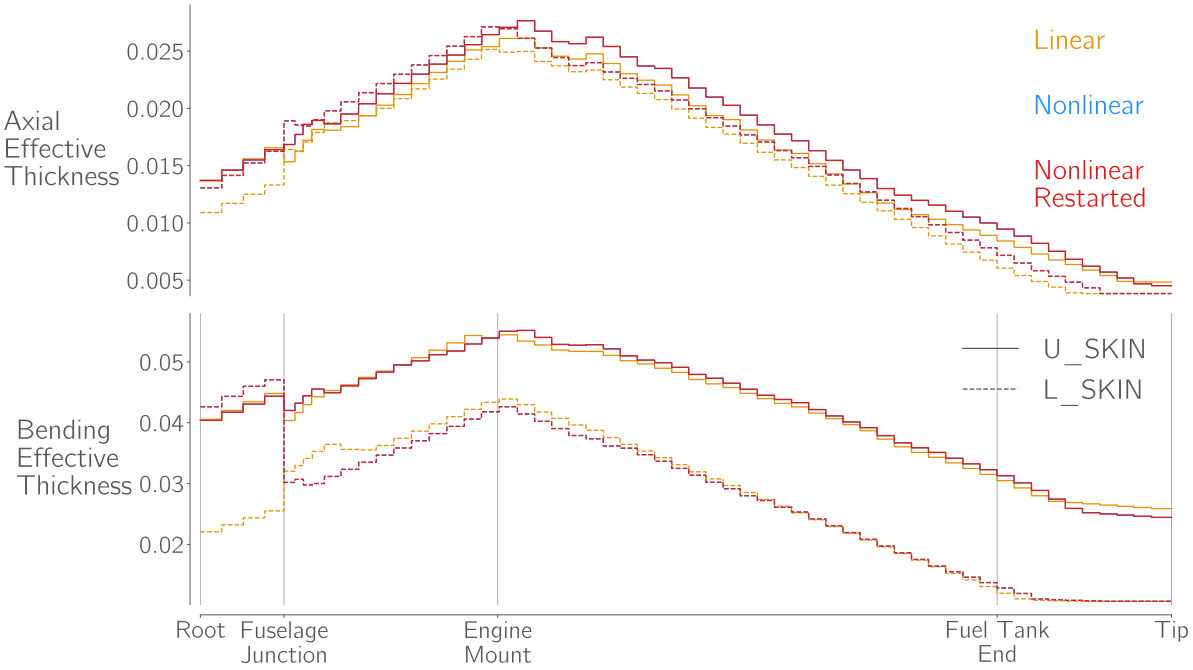

Figure 5.24: The optimisations starting from the baseline uCRM and starting from the linear optimised design converge to identical optima.

5.2. Aeroelastic Optimisation Studies

Finally, in this section I perform aeroelastic optimisation of the uCRM wingboxes. Although the initial aim of this thesis was to perform aerostructural optimisations, where the optimiser controls both the structural sizing and shape of the wing, it was discovered during the course of the thesis that the source code for computing partial derivatives of the nonlinear element residuals with respect to nodal coordinates contains fatal errors. The consequence of this is that derivatives with respect to design variables that alter the

geometry of the wingbox structure cannot be computed when using TACS' nonlinear element formulations. The optimisations presented here are therefore limited to optimising structural sizing considering coupled aeroelastic behaviour. These optimisations still provide useful insight as, unlike in the previously performed structural optimisations, when using coupled aerostructural analysis and adjoint derivatives, the optimiser is able to tune the stiffness distribution of the wingbox in order to positively influence the aerodynamic loads it experiences.

5.2.1. Methodology

As with the structural optimisations, I perform these aeroelastic optimisations on both uCRM models and with both linear and nonlinear structural formulations. The formulation of the aeroelastic optimisation problem is identical to the baseline structural optimisation formulation with the following exceptions:

- Functions of interest are computed using coupled aerostructural analysis at the 3 manoeuvre conditions described in section 4.1.2 and gradients are computed using MACH's coupled adjoint solver.
- Two additional trim constraints for each flight condition ensure that the generated lift matches the aircraft maximum take-off weight (MTOW) and that there is zero pitching moment.
- To satisfy these constraints, the angle of attack and tail rotation at each flight condition are added as design variables.
- To increase the aerostructural coupling of the optimisation problem, with the hope of accentuating any differences in linear and nonlinear optimised designs. I assume that the mass of the wingbox affects the aircraft's MTOW and consequently the lift required in each flight condition. As was done by Brooks et al. [81], I multiply the wingbox mass by 1.25 to produce an estimate of the total wing mass. I then add this mass to a reference mass which is calibrated such that the baseline uCRM wingbox design results in the original uCRM MTOW of 268×10^3 kg.

With the introduction of RANS CFD, the cost of computing each coupled aerostructural solution, and particularly each adjoint derivative, is far higher than in pure structural optimisation. I therefore take a number of steps to keep the runtime of each aeroelastic optimisation practical:

- I use the coarse FE and CFD meshes and run the three flight conditions in parallel on 144 cores each, giving a total of 432 cores.
- I use the structural designs from the baseline structural optimisations are used as the initial designs for the aeroelastic optimisations.
- Before starting each optimisation, trimmed aerostructural analyses are run to compute angle of attack and tail rotation values that correctly trim the aircraft for each flight condition.

5.2.2. Results

Unfortunately, I am unable to present results from the aeroelastic optimisations of the uCRM-13.5 in this section. Although the geometrically nonlinear aeroelastic optimisation of the uCRM-13.5 made good progress towards an optimised design, the linear aeroelastic optimisation exited significantly earlier due to numerical difficulties, not converging to an optimal design. It is therefore impossible to draw any valid conclusions about the effects of geometric nonlinearity by comparing the results of these optimisations and I therefore present only the results of the uCRM-9 optimisations in this section.

Figure 5.25 shows the optimality, feasibility and objective function histories from the two aeroelastic optimisations of the uCRM-9. Both optimisations ended with a SNOPT 40-41 exit code which indicates that SNOPT reached a point it could not improve but had not reached the specified optimality tolerance. This failure to reach the same optimality value as achieved in the structural optimisations is somewhat expected as it is not currently possible to achieve the same accuracy in the coupled adjoint gradients as is achieved in the pure structural or aerodynamic adjoints[6]. Better luck may be had re-running the optimisations with a tighter convergence tolerance on the coupled adjoint solver but this would significantly

increase the required computation time, which is already dominated by the coupled adjoint computation. Both optimisations do however achieve the desired feasibility and the structural mass shows almost no change over the second half of the optimisations which is an encouraging if not conclusive indication that the optimiser has converged to an optima.

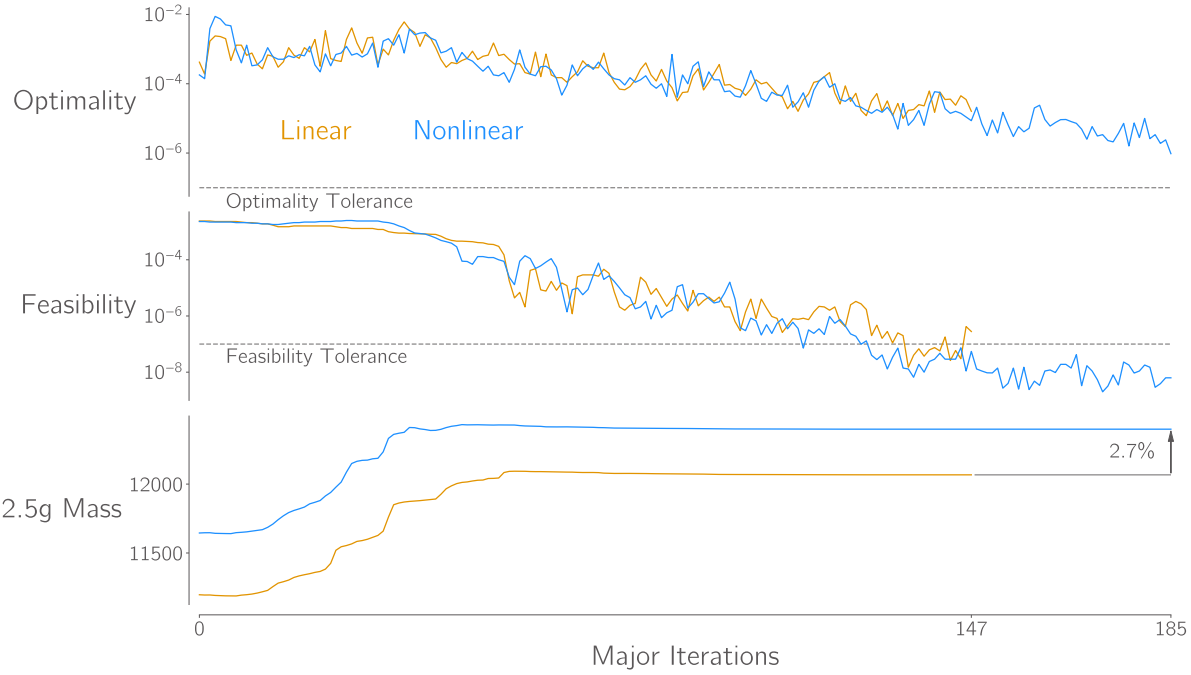

Figure 5.25: Convergence histories of the uCRM-9 aeroelastic optimisations.

The difference in mass between the linear and nonlinear aeroelastically optimised designs is 3.2%, slightly less than the 4% difference seen in the structurally optimised designs. The skin and rib sizing distributions, shown in figures 5.27 and 5.28 are very similar to those seen in the structurally optimised designs. As with those designs, the main increase in skin thickness from linear to nonlinear optimised designs is in the upper skin around the engine mount and in the centre wingbox, and the majority of the ribs are identically sized in the linear and nonlinear cases due to the previously discussed issues with the baseline structural design parameterisation.

The most striking difference between the structurally and aeroelastically optimised designs is in the sizing of the spars, shown in figure 5.29. In both the linear and nonlinear cases, the optimiser greatly reinforces the leading edge spar in order to shift the elastic axis of the wing forwards and induce more passive load alleviation under bending. Although this is a significant difference between the structurally and aeroelastically optimised designs, the spar mass is the least affected of the four component groups by the addition of nonlinearity. The lift and twist distributions in figure 5.30 show that the linear and nonlinear optimised designs achieve almost identical levels of passive load alleviation.

The fact that the geometrically nonlinear optimisation still produces the heavier design is somewhat surprising given the results of the aerostructural analysis studies, which showed that adding geometric nonlinearity decreased the bending stresses in the wing, and the fact that the above results show no significant difference in the level of passive load alleviation achieved by the two designs. The most likely cause for this unexpected result is the unphysical stress peaks present in distorted elements in the FE mesh. To demonstrate this, figure 5.31 shows the Von Mises failure criterion distributions in the upper skin of the uCRM-9 taken from the trimmed 2.5 g aerostructural analysis studies. Dispersed sparsely over the skin are quadrilateral elements which are distorted and exhibit unrealistically high stresses. Despite seeing lower stresses over the majority of the skin, it appears that the nonlinear analysis also amplifies unrealistic stress peaks and, as a result, the aggregated stress constraint value is slightly higher in the nonlinear analysis. Consequently, panels in the wingbox with these distorted elements are significantly oversized by the op-

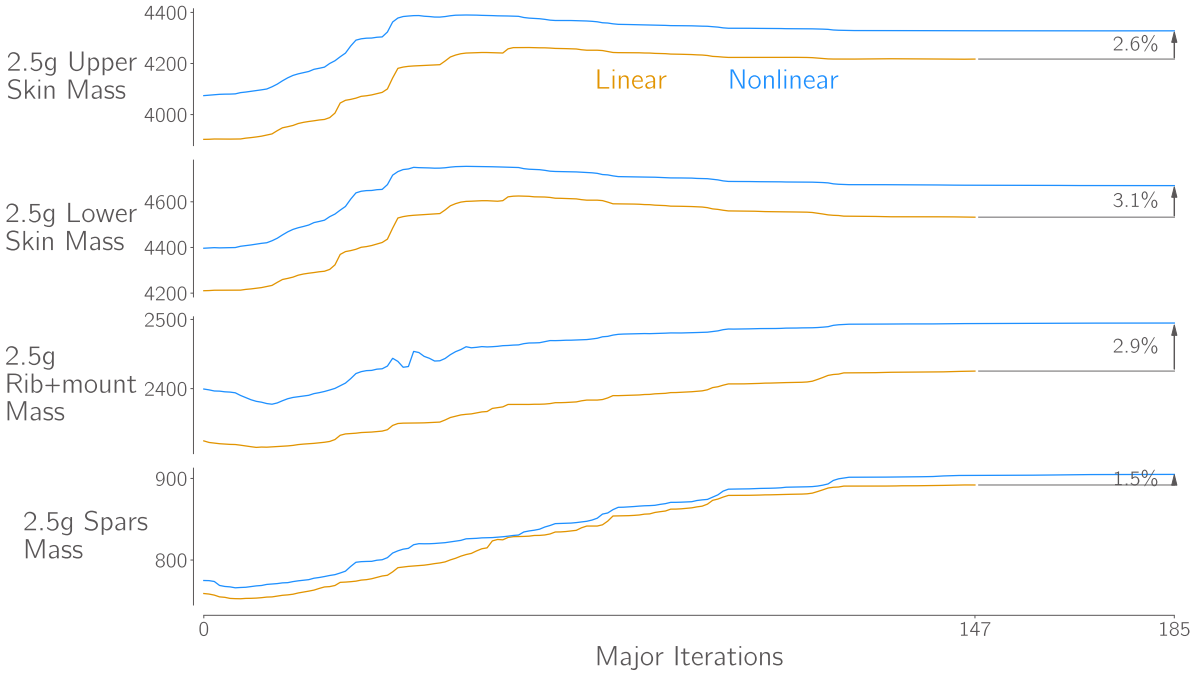

Figure 5.26: Component mass histories of the uCRM-9 aeroelastic optimisations.

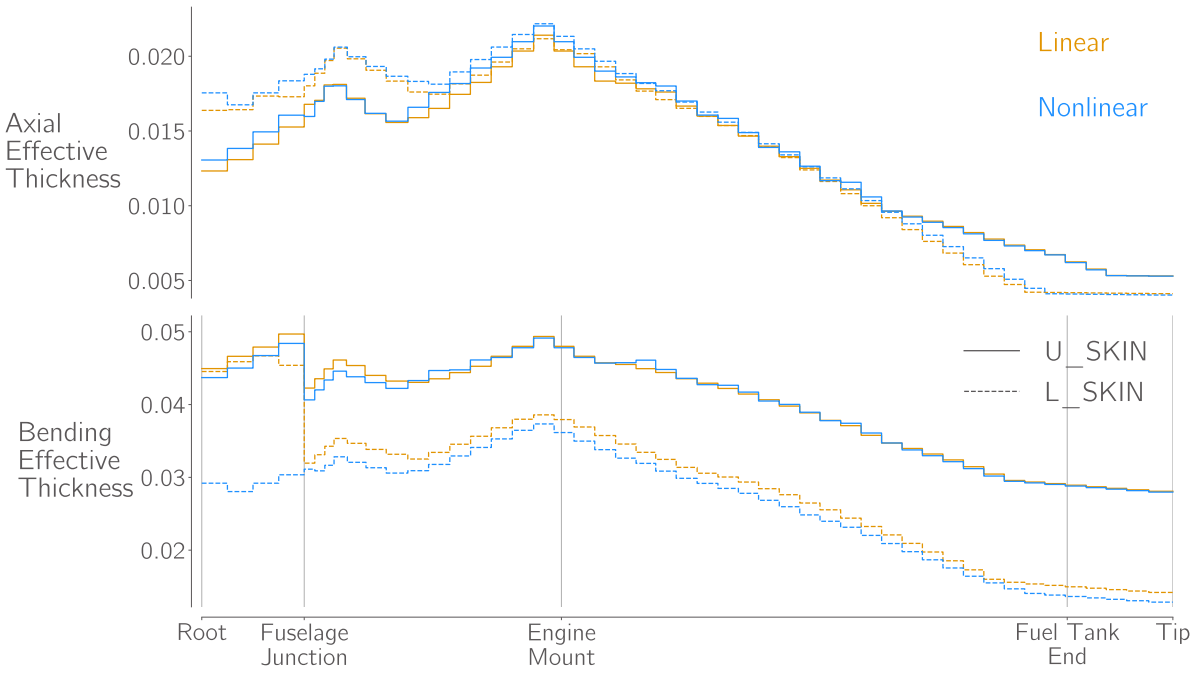

Figure 5.27: The skin sizing distributions in the linear and nonlinear aeroelastically optimised wings look very similar to those from the purely structurally optimised designs.

timiser and, if tight adjacency constraints are used, so too are adjacent panels despite their lack of stress peaks. This unfortunate oversizing will be present in both the linear and nonlinear optimised designs but, crucially, it appears that this effect is stronger in the nonlinear case, thus resulting in a heavier optimised wingbox despite generally experiencing lower bending stresses.

Another possible explanation is that the aeroelastic design space contains local optima and that each aeroelastic optimisation has converged to a different local optimum nearest it's initial structurally optimised

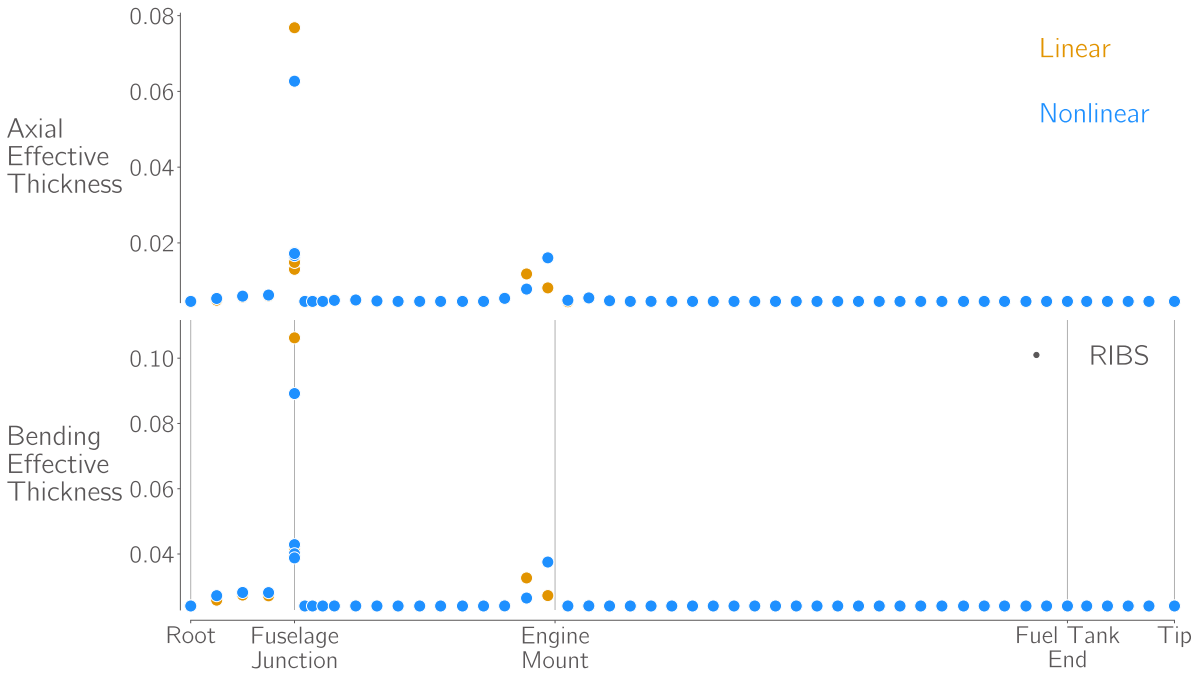

Figure 5.28: As seen in the structural optimisation studies, there is no significant difference in rib sizing between the linear and nonlinear optimised designs when using the baseline parameterisation that uses the same stiffener height and pitch on all ribs.

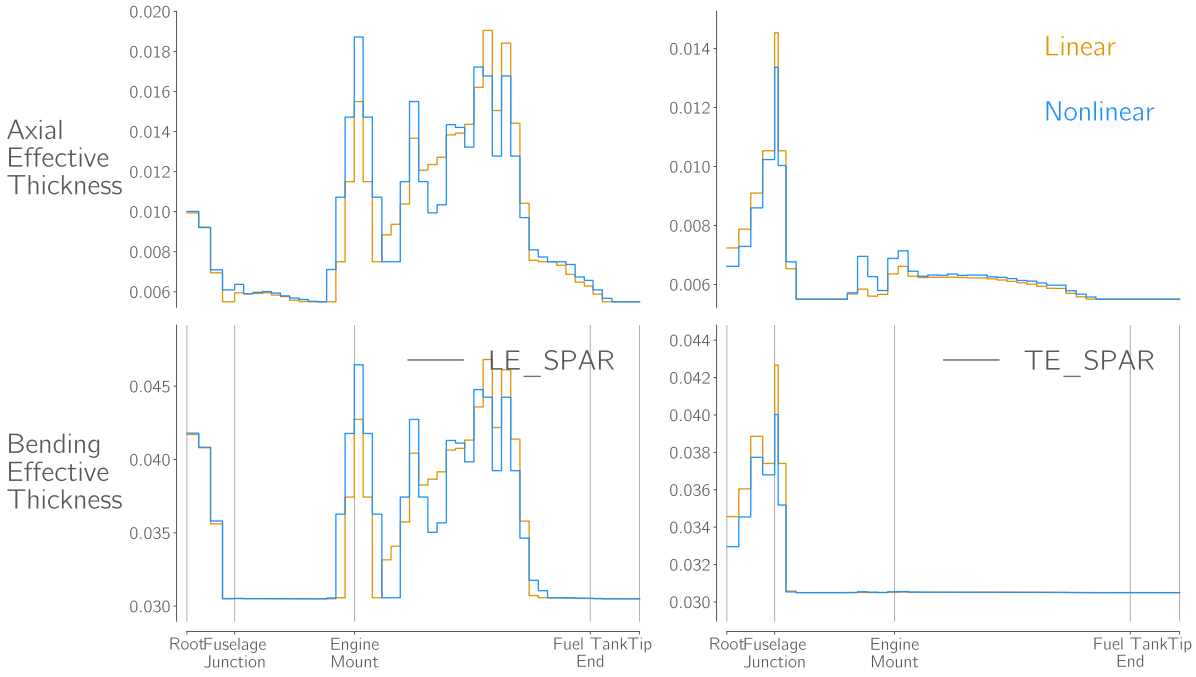

Figure 5.29: The leading edge spar is significantly reinforced in the aeroelastically optimised designs in order to increase passive load alleviation.

design. This could be investigated by repeating the linear and nonlinear optimisations starting from the same initial design but the explanation seems unlikely, given the results shown in the previous section, where the standard and bi-level structural optimisation approaches converged to the same optimum from very different initial designs.

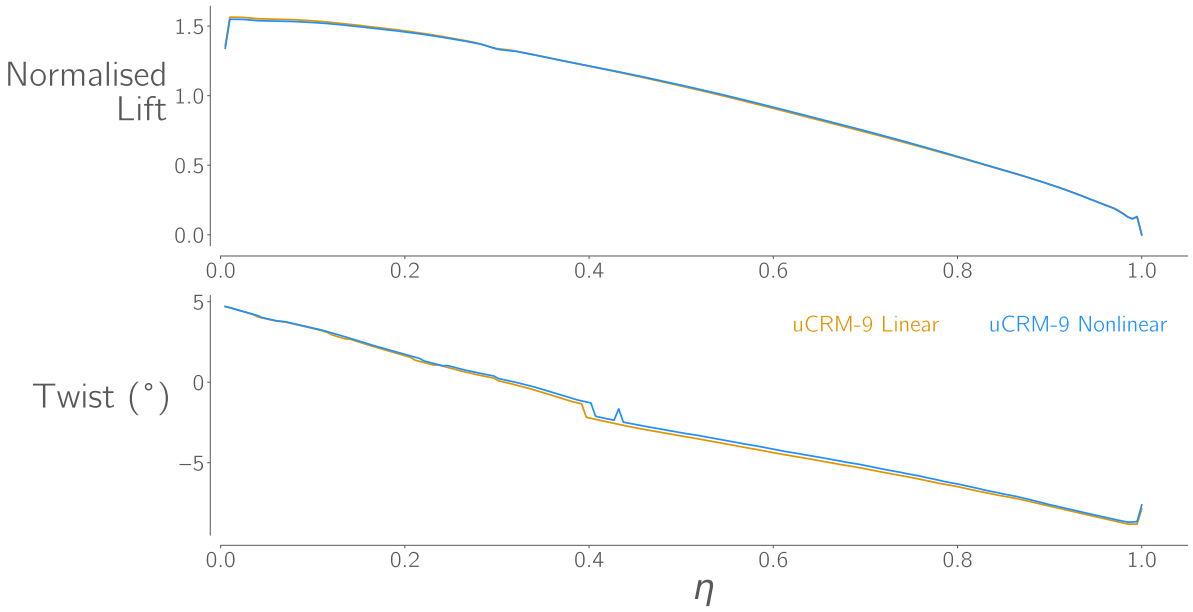

Figure 5.30: The spanwise lift and twist distributions in the two optimised uCRM-9 wings are virtually identical.

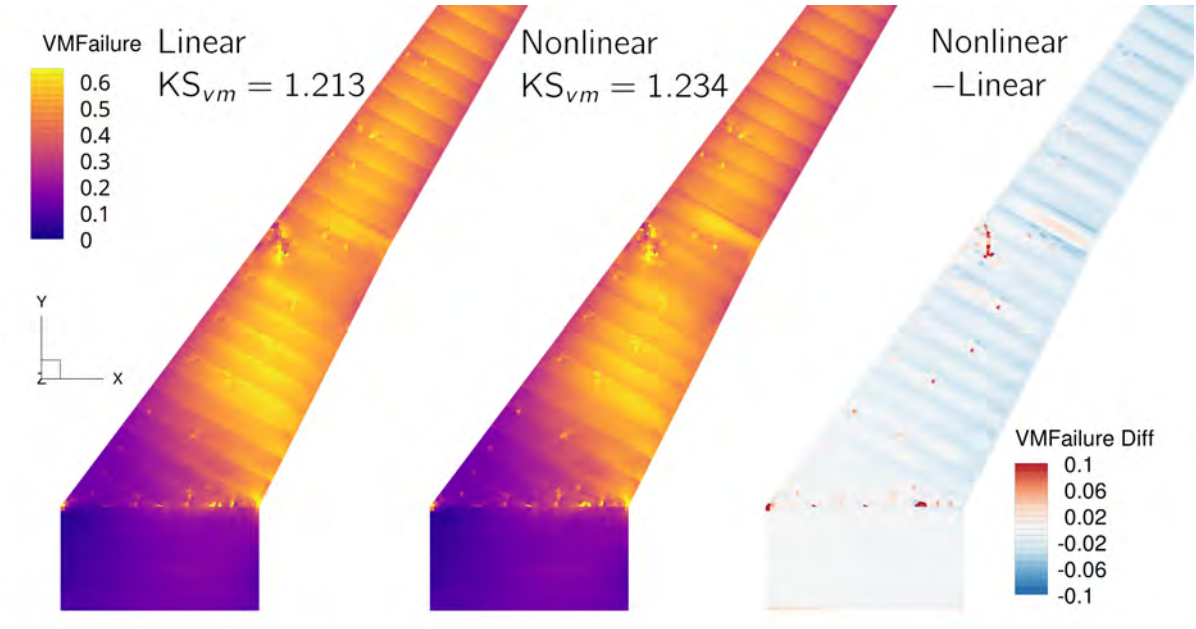

Figure 5.31: In a trimmed aerostructural analysis, even though the majority of the wingbox sees lower stresses, unrealistic stress peaks in distorted elements result in a higher failure constraint value in the nonlinear case.

Table 5.4 compares the number of optimiser iterations and function evaluations and the runtimes of the linear and nonlinear uCRM-9 aeroelastic optimisations. The geometrically nonlinear optimisation takes 26% more major iterations before exiting, but only requires 19% more runtime resulting in a 6% lower average time per major iteration. This does not indicate that the nonlinear aerostructural analysis process is any cheaper than the linear analysis but simply that in the nonlinear optimisation, fewer function evaluations were required per major iteration, this is expected as the optimiser typically requires fewer line searches in the later stages of the optimisation. The actual increase in the computational cost of each function evaluation due to the addition of the nonlinear structural formulation is only 14%, a far lower increase than seen in the analysis studies or the structural optimisations.

The reasons for this relatively small penalty associated with geometric nonlinearity during aeroelastic op-

Table 5.4: Comparison of time and number of optimiser iterations required for each aeroelastic optimisation.

Value	Linear	Nonlinear	Δ(%)	Units
Major Iterations	147	185	26	
Function Evaluations	558	581	4	
Time	53510	63592	19	s
Time per Major Iteration	364	343	-6	s
Time per Function Evaluation	96	109	14	s

timisation are threefold:

1. With the mesh sizes used here, the aerostructural solution time is still dominated by the time spent in the CFD solver.
2. Most times structural solver is called it can be restarted efficiently from the previous solution.
3. The coupled adjoint gradient computation, which dominates the computational cost of the aeroelastic optimisation, is no more expensive in the nonlinear case.

5.3. Chapter Summary

In this chapter, I performed structural and aerostructural optimisations of the undeflected common research model (uCRM)s using both linear and geometrically nonlinear structural analysis formulations in order to asses the effects of geometrically nonlinear behaviour on the optimal design of high-aspect-ratio wings (HARW).

In section 5.1, I perform structural optimisation of both uCRM wingboxes under fixed loading. The optimisations minimised mass with respect to 6-700 structural sizing variables, subject to 15 failure constraints across 3 loadcases and 7-800 linear sizing constraints. With the baseline problem formulation, using geometrically nonlinear structural analysis lead to a 6% increase in the mass of the optimised high aspect-ratio uCRM-13.5 wingbox, and a smaller but still significant 4% increase in mass in the moderate aspect-ratio uCRM-9. These mass increases occurred mostly in the skin and are consistent with the increases in bending stresses seen in the structural analysis studies presented in chapter 4. I demonstrated that, given sufficient design freedom, the optimiser was able to correctly size all the ribs in the wingbox and a significant increase in bending stiffness is seen in the ribs sized using nonlinear analysis due to the presence of Brazier loads. The mass penalty due to these loads is however relatively small. I also demonstrated that the cost of geometrically nonlinear structural optimisation is roughly an order of magnitude greater than when using linear analysis. When starting from a poor initial design, I showed that this slowdown can be reduced by up to 40% by using a bi-level approach in which an initial optimisation using linear analysis is performed to generate a good starting point for a subsequent geometrically nonlinear optimisation.

In section 5.2 I presented the results of an aeroelastic optimisation of the uCRM-9, using the same optimisation formulation as in the baseline structural optimisation but this time using coupled aerostructural analysis and gradient computation, allowing the optimiser control over the aerodynamic load distribution through aeroelastic tailoring of the structural sizing. I showed the significant differences in the sizing of the aeroelastically and structurally optimised wingboxes, with the leading edge spar being significantly reinforced in the aeroelastically optimised designs in order to increase passive load alleviation. The mass increase from linear to nonlinear aeroelastically optimised wingboxes was smaller than seen in the structural optimisations but remains inconsistent with the overall decrease in bending stresses demonstrated when comparing geometrically linear and nonlinear aerostructural analysis results in chapter 4. I posited that this unexpected mass increase is caused by unrealistic stress peaks in distorted elements within the structural mesh. These stress peaks appear to be amplified more, relative to the stresses in surrounding elements, in the nonlinear

structural analyses which causes the wingboxes optimised with nonlinear analysis to be more conservatively sized. I compared the spanwise lift and twist distributions of the two optimised designs and concluded that, at least for wings of a moderate aspect-ratio such as the uCRM-9, geometrically nonlinearity has a negligible effect on passive load alleviation. Finally, I demonstrated that including geometric nonlinearity only incurs a 14% in the computational cost of the aeroelastic optimisation.

6

Conclusions

This thesis was motivated by the current knowledge gap concerning the optimal aerostructural design of high-aspect-ratio wings (HARW) resulting from a capability gap between low-fidelity aeroelastic optimisation frameworks that include geometric nonlinearities and high-fidelity aeroelastic optimisation frameworks that cannot. The work done to date on the effect of geometric nonlinearity on the structural sizing of HARW has either involved structure-only optimisation using low-fidelity tools [22] or single analyses with high-fidelity tools which are currently too computationally expensive and lack the gradient calculation capability required for optimisation. [78]. As such, it is not clear what the effects of geometric nonlinearity are on optimal design of aircraft wings and at what point they become an important consideration.

In this thesis, I therefore aimed to do two things:

1. To develop a tool capable of performing high-fidelity aerostructural analysis and optimisation using Reynolds-averaged Navier-Stokes (RANS) computational fluid dynamics (CFD) and a geometrically nonlinear full wingbox finite element (FE) models.
2. To use the tool to investigate the consequences of geometrically nonlinear effects on the aerostructural design of modern transport aircraft.

Throughout this thesis I develop and use the MACH (MDO of Aircraft Configurations at High-fidelity) framework, developed at the University of Michigan MDO Lab, which represents the current state of the art in high-fidelity aerostructural MDO.

6.1. Contributions and Findings

To achieve the first of these aims, I implemented a Newton-Raphson based nonlinear solver in the python interface to MACH's high performance finite element library, TACS. In an attempt improve the performance of the solver I implemented multiple line search methods and an adaptive load incrementation strategy. After validating the analysis and gradient computation accuracy of the solver, I performed a large set of performance tests on two different wingbox analysis test cases of differing difficulty. The results of these tests showed that the energy based method of Matthies and Strang [102] greatly outperformed all other line search techniques but was still slower than the baseline Newton-Raphson method. The results also showed that the parameters controlling the adaptive load incrementation, the initial load increment size and the desired number of iterations per increment, strongly affect the performance of the solver. For the more challenging wingbox problem, the fastest solution times were achieved with low values for these two parameters, biasing the solver towards taking many small load steps while in the more stable test case, the opposite was true. I also introduced two methods for computing the optimal load factor from which to restart a nonlinear structural analysis from a previous solution. One based on minimising the structural residual and another novel method based on strain energy minimisation, designed to eliminate issues with the poor scaling of shell structure problems.

I then implemented some minor modifications in MACH's aerostructural Gauss-Seidel solver to improve its peformance when using the geometrically nonlinear structural formulation. The most substantial of these modifications was the addition of a load scaling ramp as originally proposed by Sanchez et al. [107] to reduce the time spent solving unrealistically highly loaded structural analyses in the early Gauss-Seidel iterations. I tested the load ramping strategy on a highly flexible transonic wing test case and showed that this ramping method can reduce time to reach a coupled solution by up to 40%.

With these solvers developed and tested, I presented I presented the results multiple studies investigating the effects of including geometric nonlinearity in structural and aerostructural analysis of modern transport aircraft with moderate and high aspect-ratio wings using both versions of the undeflected common research model (uCRM)s which represent the aspect-ratios typical of current day transport aircraft and those entering service in the next two decades. In the first study I performed linear and nonlinear structural analysis of the two uCRM wingboxes using a fixed set of loads generated by geometrically linear aerostructural analyses. I showed that, in the high aspect-ratio uCRM-13.5 wingbox, the nonlinear analysis results in greater bending stresses throughout, particularly in the upper and lower skins where increases of around 10% are observed in the Von Mises and buckling failure criteria. I demonstrated the presence of Brazier loads in the nonlinear analyses which cause a substantial increase in the compressive axial stress and buckling failure criteria in the wingbox ribs. Comparing spanwise deflection distributions, I showed that the correct modelling of large bending kinematics in the geometrically nonlinear analysis results in a significant inward and forward deflection of the wingtip, on the order of 5% of the wing semispan. I also showed that any geometrically nonlinear drag-torsion effects, if present, have little effect on the twisting deformation of the wing, which is dominated by the geometric bend-twist coupling caused by wing sweep. I proposed that the observed stress bending stress increases and the slightly increased washout in the nonlinear case are due to the nonlinear analysis correctly maintaining the bending moment caused by the lift forces on the outboard portions of the wing which are rotated inboard, a conclusion consistent with previous works demonstrating that linear analyses underpredict bending deformation when subject to follower forces. I also showed that similar differences in stress distributions are seen when performing nonlinear structural analysis on the moderate aspect-ratio uCRM-9 wing, albeit to a lesser extent. In this case, the increase in bending stresses is roughly halved, whilst the increase in the rib buckling criteria remains around 50%. The computational cost of the nonlinear analyses was greater than that of a linear analysis by a factor of 20 but depends highly on the load factor of the analysis.

Comparing the results of geometrically linear and nonlinear aerostructural analyses, both at equal angles of attack and at equal lift coefficients, I showed that steady-state aerodynamic performance values are insensitive to the levels of geometric nonlinearity seen in the uCRMs. The strongest effect of the nonlinearity on aerodynamic behaviour was seen in the pitching moment due to the span shortening and subsequent shift in the wing's centre of pressure (COP) under large bending deformations. This resulted in large relative changes in the tail rotation required to trim the aircraft but the absolute values of the changes were on the order of fractions of a degree. Changes in quantities such as the lift and drag coefficients of the aircraft were on the order of 1-4% in the highest load factor manoeuvre conditions and below 1% in a transonic cruise condition. The increase in the cost of the aerostructural analyses due to the use of nonlinear structural models was around 40%.

I then compared the wingbox stress distributions from the trimmed aerostructural analyses and showed that the nonlinear analysis resulted in 5-10% lower bending stresses in the wing skins, the opposite of the results seen when comparing structural analyses under fixed loading. To a first order approximation I proposed that this reduction in bending moment in the wingbox is due to a combination of the small reduction in lift produced by the wing and the inboard shift in the wing's COP. Buckling loads in the wing skins showed a similar decrease whilst the ribs saw a similar increase in buckling loads as seen in the structural analysis studies due to Brazier loading. Geometric nonlinearity did not appear to have a significant effect on the level of passive load alleviation achieved by either uCRM wing.

I then ran what, as far as I am aware, are the first published examples of structural and aeroelastic wingbox design optimisation using high-fidelity, geometrically nonlinear structural analysis. I performed structural optimisations to minimise the mass of the uCRM wingboxes with respect to 6-700 structural sizing variables, subject to 15 failure constraints across 3 loadcases and 7-800 linear sizing constraints.

With the baseline problem formulation, the increase in bending stresses seen in the structural analysis studies lead to a 6% increase in the mass of the optimised uCRM-13.5 wingbox, and a smaller but still significant 4% increase in mass in the uCRM-9. I also demonstrated in a subsequent optimisation that, when the optimiser is given sufficient design freedom, a significant increase in bending stiffness is seen in the ribs sized using nonlinear analysis due to the presence of Brazier loads. The mass penalty due to these loads is less than 5%. I also demonstrated that the cost of geometrically nonlinear structural optimisation is roughly an order of magnitude greater than when using linear analysis. When starting from a poor initial design, I showed that this slowdown can be reduced by up to 40% by using a bi-level approach in which an initial optimisation using linear analysis is performed to generate a good starting point for a subsequent geometrically nonlinear optimisation.

Finally, I presented the results of an aeroelastic optimisation of the uCRM-9, using the same optimisation formulation as in the baseline structural optimisation but this time using coupled aerostructural analysis and gradient computation, allowing the optimiser control over the aerodynamic load distribution through aeroelastic tailoring of the structural sizing. I showed the significant differences in the sizing of the aeroelastically and structurally optimised wingboxes, with the leading edge spar being significantly reinforced in the aeroelastically optimised designs in order to increase passive load alleviation. The mass increase from linear to nonlinear aeroelastically optimised wingboxes was smaller than seen in the structural optimisations but remains inconsistent with the overall decrease in bending stresses demonstrated when comparing geometrically linear and nonlinear aerostructural analysis comparisons. I posited that this unexpected mass increase is caused by unrealistic stress peaks in distorted elements within the structural mesh. These stress peaks appear to be amplified more, relative to the stresses in surrounding elements, in the nonlinear structural analyses which causes the wingboxes optimised with nonlinear analysis to be more conservatively sized. I compared the spanwise lift and twist distributions of the two optimised designs and concluded that, at least for wings of a moderate aspect-ratio such as the uCRM-9, geometrically nonlinearity has a negligible effect on passive load alleviation. Finally, I demonstrated that including geometric nonlinearity only incurs a 14% in the computational cost of the aeroelastic optimisation.

6.2. Answers to Research Questions

In order to assess the success of this work in achieving its original aims, I now briefly answer some of the higher level research questions posed at the start of this thesis.

1. **What is the computational cost of performing high-fidelity aerostructural analysis and optimisation with a geometrically nonlinear structural model and how can this cost be minimised?**

   (a) **How can the robustness and efficiency of a basic Newton-Raphson solver be improved for geometrically nonlinear analysis of high aspect-ratio wingboxes?** Of the line search strategies implemented in this work that were intended to improve the stability of the nonlinear structural solver, the energy based method of Matthies and Strang [102] is the top performer, being able to converge larger load increments than the pure Newton-Raphson approach in far fewer iterations than any of the residual minimisation approaches tested. The strength of this method is that it naturally accounts for the poor scaling of the translational and rotational degrees of freedom in shell problems. However, for the large deformation wingbox analyses tested in this work, foregoing any line searches and allowing the base Newton-Raphson method to take more, smaller load steps results in the fastest convergence. The BFGS updates implemented and tested in this work drastically reduced the stability of the structural solver. As a result, solutions using these approximate stiffness updates were either unobtainable or far slower than using full-order stiffness updates. Arc-length methods were not required for this work as I did not intend to explicitly model the onset of buckling in the wingbox. Buckling prediction is handled using TACS' force based buckling failure criteria. As this method provides a very conservative prediction of the wing's critical buckling modes, applying buckling failure constraints during optimisations is sufficient to avoid the optimiser producing any designs that do explicitly buckle and would not be possible to converge with the simple load incrementation method I implemented. Solving the kinds of analyses performed in this work without load incrementation

would likely be possible using a very conservative line search approach such as the minimum residual method, such an approach would however likely require 100's or 1000's of iterations to converge while using the adaptive load stepping method of Beluni and Chulya [97] allows a Newton-Raphson solver to converge the same problem in 20-30 iterations.

(b) **How can MACH's aerostructural solver be modified to work with the TACS nonlinear solver?** No modifications were required in the partitioned Gauss-Seidel aerostructural solver as this solver simply calls the implemented nonlinear structural solver, which handles the nonlinear structural solution process, including load incrementation, internally. I did however show that the aerostructural load ramping method of Sanchez et al. [107] can significantly speed up the aerostructural solver. In order to use nonlinear structural formulations with pyAeroStructure's monolithic aerostructural solver, some form of load incrementation would most likely need to be built into pyAeroStructure.

(c) **What is the overall cost penalty associated with the inclusion of structural geometric nonlinearity in structural and aerostructural analysis and optimisation?** For a standalone structural analysis, the addition of geometric nonlinearity increases the computational cost by a factor of around 20 compared to a linear analysis. However, when performing structural optimisation, the nonlinear structural solver can be restarted from the last converged solution each time it is called which reduces this slowdown factor to around 10. When performing aerostructural analysis, the nonlinear cost increase is further reduced to a factor of around 1.4 as the majority of the solution time is spent in the CFD solver. In the aeroelastic optimisations performed in this work, the computation time is dominated by the coupled adjoint solver which is no more expensive when using a nonlinear structural formulation and the penalty factor for adding nonlinearity therefore reduces to around 1.15.

2. **What effects do structural geometric nonlinearities have on the results of structural and aerostructural analysis and optimisation of HARW?**

(a) **How does geometric nonlinearity affect the results of structural and aerostructural analysis?** In structural analysis with a fixed, realistic aerodynamic load, the primary effects of geometric nonlinearity are:

i. An increased bending moment in the wingbox which leads to higher stresses in the wing skins and greater vertical deflection of the wingtip, due to the correct consideration of the effect of the rotated aerodynamic forces on the outboard portions of the wing.

ii. Compressive loads in the ribs caused by the Brazier effect.

iii. A shortening of the effective wingspan caused by geometrically correct bending kinematics.

I do not consider the increased stressing around the mounting points of the external mass loads to be a correct nonlinear effect as it is primarily caused by the method for applying these loads not being suitable for geometrically nonlinear analysis. In aerostructural analysis, the primary differences seen are caused by the shortening of the effective wingspan:

i. At the same angle of attack, the geometrically nonlinear analysis will result 1-2% less lift being produced by the deformed wing or, when trimmed to the same lift coefficient, will require a 1-2% increase in angle of attack.

ii. In swept wings, the wing's COP shifts forward, requiring less down force or more lift from the tail to trim the aircraft. As a result the wing is required to produce less lift for the same overall aircraft lift coefficient and the drag produced by the tail can change by up to 4-5 drag counts.

iii. The lack of the unrealistic stretching of the outboard portions of the wing seen with linear structural models reduces the bending moment in the wing for the same amount of lift, resulting in lower stresses in the wingbox.

Any drag torsion effects that may be present have no visible effect on the twisting of the wing

which is largely dictated by the chordwise stiffness distribution of the wingbox and the geometric bend-twist coupling caused by wing sweep. Moreover, the outboard portions of a well designed wing often produce very little drag or even a small amount of thrust, which would lead to increased washout in the wing.

(b) **How does geometric nonlinearity affect the results of structural and aerostructural optimisation?** During structural optimisation, the increased stresses resulting from the nonlinear effects described above lead to heavier optimised wingboxes. Most of this mass increase is seen in the wing skins, whose primary role is to resist bending loads in the wing. Brazier loads cause a noticeable increase in the bending stiffness of the optimised ribs but in the optimisations performed in this work, the optimiser is able to achieve this increase without much of an appreciable increase in mass, presumably by shifting material from the rib panel skins to their stiffeners. The majority of the increase in rib mass seen in the nonlinear optimisations comes from the fuselage junction rib which is subject to far more complex loading that I have not investigated in detail.

Due to issues with the computation of geometric derivatives with TACS' nonlinear elements, I was unfortunately unable to perform any true aerostructural optimisations with geometric design variables to investigate how geometric nonlinearity affect the optimal planform of HARW. I was therefore limited to using structural sizing and trim variables, performing what I refer to as aeroelastic optimisation. Optimising structural sizing considering coupled aeroelastic behaviour. Moreover, in the time available I was only able to successfully complete aeroelastic optimisations of the moderate aspect-ratio uCRM-9 which shows much weaker geometrically nonlinear behaviour than the uCRM-13.5. The results obtained showed a similar but slightly smaller increase in the mass of the uCRM-9 wingbox than was seen in the pure structural optimisations. However, based on the decrease in stress seen in the nonlinear case when comparing trimmed aerostructural analyses, it is my assertion that the geometrically nonlinear aeroelastic optimisation should result in a lower mass wingbox and that the increase in this case is caused by poor quality elements in the structural mesh as discussed above. Therefore, it is hard to draw any conclusions on the effect of geometric nonlinearity on the results of aeroelastic optimisation without further investigation.

(c) **How critical is the inclusion of geometric nonlinearity in the optimal design of HARW for future commercial transport aircraft?** Based on the results presented in this work, it appears that, although geometric nonlinearity has noticeable effects on even moderate aspect-ratio transport aircraft wings like the uCRM-9, these effects have relatively minor influence on the design of even the next generation of high aspect-ratio transport aircraft wings. This conclusion is however limited in scope as this work only considers a limited number of static aeroelastic loadcases on commercial transport aircraft wings. There are many aircraft flying today, such as sail planes or high altitude, long endurance (HALE) aircraft that have far greater aspect-ratios than the uCRM-13.5. A few of these aircraft, such as NASA's Helios and Facebook's Aquila protoypes, have famously suffered catastrophic failures due to geometrically nonlinear effects not considered during their design Additionally, other phenomena that are important for HARW design may be affected more critically by geometric nonlinearity, including control surface effectiveness, flutter and flight dynamics.

6.3. Recommendations for Future work

There are many avenues for future work improving and applying the geometrically nonlinear capabilities introduced to MACH in this work, including:

Quantifying the effect of poor quality structural meshes: In order to substantiate some of the claims I made about the effect of the distorted elements in the current uCRM structural meshes, the wingbox geometries should be remeshed, taking great care to avoid such distorted elements, before re-running the optimisations performed in this thesis.

Making MACH fully geometrically nonlinear: As I have mentioned already in this work, both RLT load-displacement transfer scheme and the method used for applying external mass loads are both still not valid under large deformations. Upgrading these methods to be geometrically exact may increase the effect of geometric nonlinearity on the kind of optimisations performed in this work. During the course of this thesis, I successfully implemented a geometrically nonlinear version of the RLT displacement transfer method but extending this to the load transfer scheme and moreover the relevant gradient computation routines is a more challenging implementation task.

Integrating nonlinear structural formulations with MACH's montolithic aerostructural solver: In this work I used only the partitioned Gauss-Seidel solver due to the ease of including nonlinear structural analysis to the solver. As a next step, efforts should be made to modify MACH's monolithic coupled Newton-Krylov aerostructural solver to work with nonlinear structural formulations due to the potential performance benefits. As mentioned above, this may require integrating some form of load incrementation into the aerostructural solver.

Enabling true aerostructural optimisation: The original end goal of this thesis was to perform simultaneous optimisation of the uCRM's structural sizing and wing shape with geometrically nonlinear aerostructural analysis and in doing so, gain an understanding on the effect of geometric nonlinearities on the optimal aspect-ratio of the next generation of transport aircraft. To enable this, the geometric derivative errors in TACS' current nonlinear elements must be overcome. This could be done either by fixing the current errors in the MITC shell elements used in MACH, or by integrating the newer version of TACS into MACH and using it's MITC9 shell elements which do not suffer from this error.

Aerostructural optimisation of extermely flexible wings: Previously, the vast majority of published work using MACH has focused on optimising commercial transport aircaft. With geometrically nonlinear capabilities added it would be interesting to apply MACH to the aerostructural design of an extremely flexible HALE aircraft. This may pose a significant challenge for MACH as the ADflow CFD code is specialised for supersonic and transonic flows and thus loses much of it's robustness and efficiency when simulating subsonic flows.

Validation of nonlinear aerostructural solver: To date, MACH's aerostructural solver has not been validated against any aeroelastic benchmark problems. An excellent candidate for such a validation case is the highly flexible Pazy wing, designed for the large deflection working group of the aeroelastic prediction workshop, which is currently undergoing a wind tunnel test campaign.

Bibliography

[1] Castellani, M., Cooper, J. E., and Lemmens, Y., "Nonlinear static aeroelasticity of high-aspect-ratio-wing aircraft by finite element and multibody methods," *Journal of Aircraft*, Vol. 54, No. 2, 2017, pp. 548–560. doi:10.2514/1.C033825.

[2] Brooks, T. R., Martins, J. R. R. A., and Kennedy, G. J., "High-fidelity Aerostructural Optimization of Tow-steered Composite Wings," *Journal of Fluids and Structures*, Vol. 88, 2019, pp. 122–147. doi:10.1016/j.jfluidstructs.2019.04.005.

[3] Burdette, D. A., and Martins, J. R. R. A., "Impact of Morphing Trailing Edge on Mission Performance for the Common Research Model," *Journal of Aircraft*, Vol. 56, No. 1, 2019, pp. 369–384. doi: 10.2514/1.C034967.

[4] Kennedy, G. J., Kenway, G. K. W., and Martins, J. R. R. A., "High Aspect Ratio Wing Design: Optimal Aerostructural Tradeoffs for the Next Generation of Materials," *Proceedings of the AIAA Science and Technology Forum and Exposition (SciTech)*, National Harbor, MD, 2014. doi:10.2514/6.2014-0596.

[5] Kenway, G. W. K., and Martins, J. R. R. A., "High-fidelity aerostructural optimization considering buffet onset," *Proceedings of the 16th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*, Dallas, TX, 2015. AIAA 2015-2790.

[6] Kenway, G. K. W., Kennedy, G. J., and Martins, J. R. R. A., "Scalable Parallel Approach for High-Fidelity Steady-State Aeroelastic Analysis and Derivative Computations," *AIAA Journal*, Vol. 52, No. 5, 2014, pp. 935–951. doi:10.2514/1.J052255.

[7] Dowell, E., Edwards, J., and Strganac, T. W., "Nonlinear aeroelasticity," *Collection of Technical Papers - AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics and Materials Conference*, Vol. 5, No. 5, 2003, pp. 3824–3847. doi:10.2514/2.6876.

[8] Patil, M. J., and Hodges, D. H., "On the importance of aerodynamic and structural geometrical nonlinearities in aeroelastic behavior of high-aspect-ratio wings," *Journal of Fluids and Structures*, Vol. 19, No. 7, 2004, pp. 905–915. doi:10.1016/j.jfluidstructs.2004.04.012.

[9] Garcia, J. A., "Numerical investigation of nonlinear aeroelastic effects on flexible high-aspect-ratio wings," *Journal of Aircraft*, Vol. 42, No. 4, 2005, pp. 1025–1036. doi:10.2514/1.6544.

[10] Su, W., "Coupled Nonlinear Aeroelasticity and Flight Dynamics of Fully Flexible Aircraft," Phd, University of Michigan, 2008. doi:10.1007/s13398-014-0173-7.2.

[11] Cavagna, L., Ricci, S., and Riccobene, L., "Structural sizing, aeroelastic analysis, and optimization in aircraft conceptual design," *Journal of Aircraft*, Vol. 48, No. 6, 2011, pp. 1840–1855. doi: 10.2514/1.C031072.

[12] Harmin, M. Y., and Cooper, J. E., "Aeroelastic behaviour of a wing including geometric nonlinearities," *Aeronautical Journal*, Vol. 115, No. 1174, 2011, pp. 767–777. doi:10.1017/S0001924000006515.

[13] Palacios, R., Wang, Y., Wynn, A., and Karpel, M., "Condensation of large finite-element models for wing load analysis with geometrically-nonlinear effects," *IFASD 2013 - International Forum on Aeroelasticity and Structural Dynamics*, 2013, pp. 1–25.

[14] Bartels, R. E., Scott, R. C., Allen, T. J., and Sexton, B. W., "Aeroelastic analysis of SUGAR truss-braced wing wind-tunnel model using FUN3D and a nonlinear structural model," *56th AIAA/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference*, 2015. doi:

10.2514/6.2015-1174.

[15] Howcroft, C., Cook, R., Calderon, D., Lambert, L., Castellani, M., Cooper, J. E., Lowenberg, M. H., and Neild, S., "Aeroelastic modelling of highly flexible wings," *15th Dynamics Specialists Conference*, AIAA, San Diego, 2016. doi:10.2514/6.2016-1798.

[16] Castellani, M., Cooper, J. E., and Lemmens, Y., "Nonlinear static aeroelasticity of high-aspect-ratio-wing aircraft by finite element and multibody methods," *Journal of Aircraft*, Vol. 54, No. 2, 2017, pp. 548–560. doi:10.2514/1.C033825.

[17] Ritter, M., Teixeira, P. C., and Cesnik, C. E., "Comparison of nonlinear aeroelastic methods for maneuver simulation of very flexible aircraft," *AIAA/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference, 2018*, 2018. doi:10.2514/6.2018-1953.

[18] Riso, C., Di Vincenzo, F. G., Ritter, M., Cesnik, C. E., and Mastroddi, F., "Nonlinear aeroelastic trim of very flexible aircraft described by detailed models," *Journal of Aircraft*, Vol. 55, No. 6, 2018, pp. 2338–2346. doi:10.2514/1.C034787.

[19] Lupp, C. A., and Cesnik, C. E. S., "A Gradient-Based Flutter Constraint Including Geometrically Nonlinear Deformations," *2019 AIAA/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference*, AIAA, San Diego, California, 2019.

[20] Medeiros, R. R., Cesnik, C. E. S., and Coetzee, E. B., "Computational Aeroelasticity Using Modal-Based Structural Nonlinear Analysis," *AIAA Journal*, 2019, pp. 1–10. doi:10.2514/1.j058593.

[21] Werter, N. P., and De Breuker, R., "A novel dynamic aeroelastic framework for aeroelastic tailoring and structural optimisation," *Composite Structures*, Vol. 158, 2016, pp. 369–386. doi:10.1016/j.compstruct.2016.09.044, URL http://dx.doi.org/10.1016/j.compstruct.2016.09.044.

[22] Calderon, D. E., Cooper, J. E., Lowenberg, M., Neild, S. A., and Coetzee, E. B., "Sizing High-Aspect-Ratio Wings with a Geometrically Nonlinear Beam Model," *Journal of Aircraft*, Vol. 56, No. 4, 2019, pp. 1455–1470. doi:10.2514/1.c035296.

[23] Krüger, W. R., Dillinger, J., De Breuker, R., and Haydn, K., "Investigations of passive wing technologies for load reduction," *CEAS Aeronautical Journal*, Vol. 10, No. 4, 2019, pp. 977–993. doi:10.1007/s13272-019-00393-2, URL https://doi.org/10.1007/s13272-019-00393-2.

[24] Ashley, H., "On Making Things the Best—Aeronautical Uses of Optimization," *Journal of Aircraft*, Vol. 19, No. 1, 1982, pp. 5–28. doi:10.2514/3.57350.

[25] Triplett, W. E., "Aeroelastic tailoring studies in fighter aircraft design," *Journal of Aircraft*, Vol. 17, No. 7, 1980, pp. 508–513. doi:10.2514/3.57932.

[26] Love, M., and Bohlman, J., "Aeroelastic Tailoring and Integrated Wing Design," Tech. Rep. 89N25167, NASA, 1989. URL https://ntrs.nasa.gov/archive/nasa/casi.ntrs.nasa.gov/19890015796.pdf.

[27] Haftka, R. T., "Optimization of Flexible Wing Structures Subject to Strength and Induced Drag Constraints," *AIAA Journal*, Vol. 15, No. 8, 1977, pp. 1101–1106. doi:10.2514/3.7400.

[28] Grossman, B., Gurdal, Z., Strauch, G. J., Eppard, W. M., and Haftka, R. T., "Integrated Aerodynamic/Structural Design of a Sailplane Wing," *Journal of Aircraft*, Vol. 25, No. 9, 1988, pp. 855–860. doi:10.2514/3.45670.

[29] Grossman, B., Haftka, R. T., Kao, P.-J., Polen, D. M., and Rais-Rohani, M., "Integrated Aerodynamic-Structural Design of a Transport Wing," *Journal of Aircraft*, Vol. 27, No. 12, 1990, pp. 1050–1056. doi:10.2514/3.45980.

[30] Sobieszczanski-Sobieski, J., "Sensitivity analysis and multidisciplinary optimization for aircraftdesign-recent advances and results," *Journal of Aircraft*, Vol. 27, No. 12, 1990, pp. 993–1001. doi:10.2514/3.45973.

[31] Adelman, H. M., and Haftka, R. T., "Sensitivity Analysis of Discrete Structural Systems," *AIAA Journal*, Vol. 24, No. 5, 1986, pp. 823–832. doi:10.2514/3.48671.

[32] Maute, K., Lesoinne, M., and Farhat, C., "Optimization of Aeroelastic Systems Using Coupled Analytical Sensitivities," *Proceedings of the 38th AIAA Aerospace Sciences Meeting*, Reno, NV, 2000. AIAA 2000-0560.

[33] Maute, K., Nikbay, M., and Farhat, C., "Coupled Analytical Sensitivity Analysis and Optimization of Three-Dimensional Nonlinear Aeroelastic Systems," *AIAA Journal*, Vol. 39, No. 11, 2001, pp. 2051–2061.

[34] Maute, K., Nikbay, M., and Farhat, C., "Sensitivity Analysis and Design Optimization of Three-Dimensional Non-Linear Aeroelastic Systems by the Adjoint Method," *International Journal for Numerical Methods in Engineering*, Vol. 56, No. 6, 2003, pp. 911–933.

[35] Barcelos, M., Bavestrello, H., and Maute, K., "A Schur–Newton–Krylov solver for steady-state aeroelastic analysis and design sensitivity analysis," *Computer Methods in Applied Mechanics and Engineering*, Vol. 195, 2006, pp. 2050—2069.

[36] Barcelos, M., and Maute, K., "Aeroelastic design optimization for laminar and turbulent flows," *Computer Methods in Applied Mechanics and Engineering*, Vol. 197, 2008, pp. 1813–1832. doi:10.1016/j.cma.2007.03.009.

[37] Martins, J. R. R. A., Alonso, J. J., and Reuther, J. J., "Aero-Structural Wing Design Optimization Using High-Fidelity Sensitivity Analysis," *Proceedings of the CEAS Conference on Multidisciplinary Aircraft Design and Optimization*, edited by H. Hölinger, Köln, Germany, 2001, pp. 211–226.

[38] Martins, J. R. R. A., Alonso, J. J., and Reuther, J. J., "A Coupled-Adjoint Sensitivity Analysis Method for High-Fidelity Aero-Structural Design," *Optimization and Engineering*, Vol. 6, No. 1, 2005, pp. 33–62. doi:10.1023/B:OPTE.0000048536.47956.62.

[39] Martins, J. R. R. A., Alonso, J. J., and Reuther, J. J., "High-Fidelity Aerostructural Design Optimization of a Supersonic Business Jet," *Journal of Aircraft*, Vol. 41, No. 3, 2004, pp. 523–530. doi:10.2514/1.11478.

[40] Mader, C. A., Martins, J. R. R. A., Alonso, J. J., and van der Weide, E., "ADjoint: An Approach for the Rapid Development of Discrete Adjoint Solvers," *AIAA Journal*, Vol. 46, No. 4, 2008, pp. 863–873. doi:10.2514/1.29123.

[41] Lyu, Z., Kenway, G. K., Paige, C., and Martins, J. R. R. A., "Automatic Differentiation Adjoint of the Reynolds-Averaged Navier–Stokes Equations with a Turbulence Model," *21st AIAA Computational Fluid Dynamics Conference*, San Diego, CA, 2013. doi:10.2514/6.2013-2581.

[42] Kennedy, G. J., and Martins, J. R. R. A., "A Parallel Finite-Element Framework for Large-Scale Gradient-Based Design Optimization of High-Performance Structures," *Finite Elements in Analysis and Design*, Vol. 87, 2014, pp. 56–73. doi:10.1016/j.finel.2014.04.011.

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

[44] Kennedy, G. J., and Martins, J. R. R. A., "A parallel aerostructural optimization framework for aircraft design studies," *Structural and Multidisciplinary Optimization*, Vol. 50, No. 6, 2014, pp. 1079–1101. doi:10.1007/s00158-014-1108-9.

[45] Burdette, D., Kenway, G. K. W., Lyu, Z., and Martins, J. R. R. A., "Aerostructural Design Optimization of an Adaptive Morphing Trailing Edge Wing," *Proceedings of the AIAA Science and Technology Forum and Exposition (SciTech)*, Kissimmee, FL, 2015. doi:10.2514/6.2016-1294.

[46] Burdette, D. A., Kenway, G. K., and Martins, J. R. R. A., "Performance Evaluation of a Morphing Trailing Edge Using Multipoint Aerostructural Design Optimization," *57th AIAA Structures,*

Structural Dynamics, and Materials Conference, AIAA, 2016. doi:10.2514/6.2016-0159.

[47] Burdette, D. A., Kenway, G. K. W., and Martins, J. R. R. A., "Aerostructural design optimization of a continuous morphing trailing edge aircraft for improved mission performance," *17th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*, Washington, D.C., 2016. doi:10.2514/6.2016-3209.

[48] Burdette, D. A., "High-Fidelity Aerostructural Design Optimization of Transport Aircraft with Continuous Morphing Trailing Edge Technology," Ph.D. thesis, University of Michigan, 2017.

[49] Burdette, D. A., and Martins, J. R. R. A., "Design of a Transonic Wing with an Adaptive Morphing Trailing Edge via Aerostructural Optimization," *Aerospace Science and Technology*, Vol. 81, 2018, pp. 192–203. doi:10.1016/j.ast.2018.08.004.

[50] Brooks, T. R., Hwang, J. T., Kennedy, G. J., and Martins, J. R. R. A., "High-fidelity structural optimization of a tow-steered composite wing," *Proceedings of the 11th World Congress on Structural and Multidisciplinary Optimization*, Sydney, Australia, 2015.

[51] Brooks, T. R., Kennedy, G. J., and Martins, J. R. R. A., "High-fidelity Aerostructural Optimization of a High Aspect Ratio Tow-steered Wing," *57th AIAA/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference*, American Institute of Aeronautics and Astronautics, 2016. doi:10.2514/6.2016-1179.

[52] Brooks, T. R., Kennedy, G. J., and Martins, J. R. R. A., "High-fidelity Multipoint Aerostructural Optimization of a High Aspect Ratio Tow-steered Composite Wing," *Proceedings of the 58th AIAA/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference, AIAA SciTech Forum*, Grapevine, TX, 2017. doi:10.2514/6.2017-1350.

[53] Brooks, T. R., "Design Optimization of Flexible Aircraft Wings Using Tow-steered Composites," Ph.D. thesis, University of Michigan, Ann Arbor, MI, 2018.

[54] Brooks, T. R., and Martins, J. R. R. A., "On Manufacturing Constraints for Tow-steered Composite Design Optimization," *Composite Structures*, Vol. 204, 2018, pp. 548–559. doi:10.1016/j.compstruct.2018.07.100.

[55] Jasa, J. P., Hwang, J. T., and Martins, J. R. R. A., "Open-source coupled aerostructural optimization using Python," *Structural and Multidisciplinary Optimization*, Vol. 57, No. 4, 2018, pp. 1815–1827. doi:10.1007/s00158-018-1912-8.

[56] Chauhan, S. S., and Martins, J. R. R. A., "Low-Fidelity Aerostructural Optimization of Aircraft Wings with a Simplified Wingbox Model Using OpenAeroStruct," *Proceedings of the 6th International Conference on Engineering Optimization, EngOpt 2018*, Springer, Lisbon, Portugal, 2018, pp. 418–431. doi:10.1007/978-3-319-97773-7_38.

[57] Elham, A., and van Tooren, M. J., "Coupled adjoint aerostructural wing optimization using quasi-three-dimensional aerodynamic analysis," *Structural and Multidisciplinary Optimization*, Vol. 54, No. 4, 2016, pp. 889–906. doi:10.1007/s00158-016-1447-9.

[58] Jasa, J. P., Chauhan, S. S., Gray, J. S., and Martins, J. R. R. A., "How Certain Physical Considerations Impact Aerostructural Wing Optimization," *AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*, Dallas, TX, 2019. doi:10.2514/6.2019-3242.

[59] Venkatamaran, S., and Haftka, R. T., "Structural optimization complexity: what has Moore's law done for us?" *Structural and Multidisciplinary Optimization*, Vol. 28, 2004, pp. 375–387. doi:10.1007/s00158-004-0415-y.

[60] Rajpal, D., Gillebaart, E., and De Breuker, R., "Preliminary aeroelastic design of composite wings subjected to critical gust loads," *Aerospace Science and Technology*, Vol. 85, 2019, pp. 96–112. doi:10.1016/j.ast.2018.11.051, URL https://doi.org/10.1016/j.ast.2018.11.051.

[61] Stanford, B., "Aeroservoelastic Optimization of a Transport Aircraft Wingbox under Stochastic Gust

Constraints," *Journal of Aeroelasticity and Structural Dynamics*, Vol. 6, No. 1, 2018, pp. 21–41. doi:10.3293/asdj.2018.48.

[62] Engelsen, F., and Livne, E., "Mode Acceleration Based Random Gust Stresses in Aeroservoelastic Optimization," *Journal of Aircraft*, Vol. 41, No. 2, 2004, pp. 335–347. doi:10.2514/1.9329.

[63] Rajpal, D., Kassapoglou, C., and De Breuker, R., "Aeroelastic optimization of composite wings including fatigue loading requirements," *Composite Structures*, Vol. 227, No. August, 2019, p. 111248. doi:10.1016/j.compstruct.2019.111248, URL `https://doi.org/10.1016/j.compstruct.2019.111248`.

[64] Stanford, B. K., "Aeroelastic Wingbox Stiffener Topology Optimization," *Journal of Aircraft*, Vol. 55, No. 3, 2017, pp. 1244–1251. doi:10.2514/1.C034653.

[65] Stanford, B. K., and Dunning, P. D., "Optimal Topology of Aircraft Rib and Spar Structures Under Aeroelastic Loads," *Journal of Aircraft*, Vol. 52, No. 4, 2014, pp. 1298–1311. doi:10.2514/1.C032913.

[66] Stanford, B. K., "Shape , Sizing , and Topology Design of a Wingbox under Aeroelastic Constraints," *AIAA Aviation 2020 Forum*, 2020. doi:10.2514/6.2020-3147.

[67] Stanford, B. K., Jutte, C. V., and Coker, C. A., "Aeroelastic Sizing and Layout Design of a Wingbox Through Nested Optimization," *AIAA Journal*, Vol. 57, No. 2, 2018, pp. 848–857. doi:10.2514/1.j057428.

[68] Lancelot, P., and De Breuker, R., "Transonic Flight and Movable Load Modelling for Wing-Box Preliminary Sizing," *18th International Forum on Aeroelasticity and Structural Dynamics*, , No. June, 2019, pp. 28–30.

[69] Natella, M., and De Breuker, R., "The effects of a full-aircraft aerodynamic model on the design of a tailored composite wing," *CEAS Aeronautical Journal*, Vol. 10, No. 4, 2019, pp. 995–1014. doi:10.1007/s13272-019-00366-5, URL `http://dx.doi.org/10.1007/s13272-019-00366-5`.

[70] Bordogna, M. T., Lancelot, P., Bettebghor, D., and De Breuker, R., "Aeroelastic tailoring for static and dynamic loads with blending constraints," *17th International Forum on Aeroelasticity and Structural Dynamics, IFASD 2017*, Vol. 2017-June, No. June, 2017.

[71] Dillinger, J. K., Abdalla, M. M., Meddaikar, Y. M., and Klimmek, T., "Static aeroelastic stiffness optimization of a forward swept composite wing with CFD-corrected aero loads," *CEAS Aeronautical Journal*, Vol. 10, No. 4, 2019, pp. 1015–1032. doi:10.1007/s13272-019-00397-y, URL `https://doi.org/10.1007/s13272-019-00397-y`.

[72] Jovanov, K., "High-Fidelity Load and Gradient Corrections for Static Aeroelastic Tailoring of Composite Wings," Phd, TU Delft, 2019. doi:10.4233/uuid, URL `https://doi.org/10.4233/uuid:14b55d5e-586a-4641-8990-55a397674db8`.

[73] Brazier, L. G., "On the Flexure of Thin Cylindrical Shells and Other "Thin" Sections," *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, Vol. 116, No. 773, 1927, pp. 104–114. doi:10.1098/rspa.1927.0125.

[74] Stanford, B. K., and Dunning, P. D., "Optimal Topology of Aircraft Rib and Spar Structures Under Aeroelastic Loads," *Journal of Aircraft*, Vol. 52, No. 4, 2015, pp. 1298–1311. doi:10.2514/1.C032913, URL `http://arc.aiaa.org/doi/10.2514/1.C032913`.

[75] Patil, M. J., Hodges, D. H., and S. Cesnik, C. E., "Nonlinear Aeroelasticity and Flight Dynamics of High-Altitude Long-Endurance Aircraft," *Journal of Aircraft*, Vol. 38, No. 1, 2001, pp. 88–94. doi:10.2514/2.2738, URL `https://doi.org/10.2514/2.2738`.

[76] Ritter, M., and Cesnik, C. E., "Large deformation modeling of a beam type structure and a 3D wingbox using an enhanced modal approach," *57th AIAA/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference*, , No. January, 2016, pp. 1–20. doi:10.2514/6.2016-1708.

[77] Smith, M. J., Patil, M. J., and Hodges, D. H., "CFD-based analysis of nonlinear aeroelastic behavior of high-aspect ratio wings," *19th AIAA Applied Aerodynamics Conference*, , No. c, 2001. doi: 10.2514/6.2001-1582.

[78] Verri, A. A., Bussamra, F. L., de Morais, K. C., Becker, G. G., Cesnik, C. E., Luque Filho, G. B., and de Oliveira, L. C., "Static loads evaluation in a flexible aircraft using high-fidelity fluidstructure iteration tool (E2-FSI): extended version," *Journal of the Brazilian Society of Mechanical Sciences and Engineering*, Vol. 42, No. 1, 2020. doi:10.1007/s40430-019-2154-4.

[79] Calderon, D. E., Cooper, J. E., Lowenberg, M. H., and Neild, S. A., "On the Effect of Including Geometric Nonlinearity in the Sizing of a Wing," *AIAA/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference, 2018*, Kissimmee, 2018. doi:10.2514/6.2018-1680.

[80] Stanford, B., Wieseman, C. D., and Jutte, C., "Aeroelastic Tailoring of Transport Wings Including Transonic Flutter Constraints," *56th AIAA/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference*, Kissimmee, FL, 2015. doi:10.2514/6.2015-1127.

[81] Brooks, T. R., Kenway, G. K. W., and Martins, J. R. R. A., "Benchmark Aerostructural Models for the Study of Transonic Aircraft Wings," *AIAA Journal*, Vol. 56, No. 7, 2018, pp. 2840–2855. doi:10.2514/1.J056603.

[82] Mader, C. A., Kenway, G. K. W., Yildirim, A., and Martins, J. R. R. A., "ADflow—An open-source computational fluid dynamics solver for aerodynamic and multidisciplinary optimization," *Journal of Aerospace Information Systems*, 2020. doi:10.2514/1.I010796.

[83] Yildirim, A., Kenway, G. K. W., Mader, C. A., and Martins, J. R. R. A., "A Jacobian-free approximate Newton–Krylov startup strategy for RANS simulations," *Journal of Computational Physics*, Vol. 397, 2019, p. 108741. doi:10.1016/j.jcp.2019.06.018.

[84] Kenway, G. K. W., Mader, C. A., He, P., and Martins, J. R. R. A., "Effective Adjoint Approaches for Computational Fluid Dynamics," *Progress in Aerospace Sciences*, Vol. 110, 2019, p. 100542. doi:10.1016/j.paerosci.2019.05.002.

[85] Kennedy, G. J., "Aerostructural analysis and design optimization of composite aircraft," Ph.D. thesis, University of Toronto, Toronto, ON, Canada, December 2012.

[86] Martins, J. R. R. A., and Poon, N. M. K., "On Structural Optimization Using Constraint Aggregation," *Proceedings of the 6th World Congress on Structural and Multidisciplinary Optimization*, Rio de Janeiro, Brazil, 2005.

[87] Lambe, A. B., and Martins, J. R. R. A., "Matrix-free aerostructural optimization of aircraft wings," *Structural and Multidisciplinary Optimization*, Vol. 53, No. 3, 2016, pp. 589–603. doi:10.1007/s00158-015-1349-2.

[88] Iorga, L., Malmedy, V., Stodieck, O., and Loxham, J., "Preliminary sizing optimisation of aircraft structures: industrial challenges and practices," *6th Aircraft Structural Design Conference*, Bristol, UK, 2018.

[89] Bushnell, D., "Buckling of shellspitfall for designers," *AIAA Journal*, Vol. 19, No. 9, 1981, pp. 1183–1226. doi:10.2514/3.60058.

[90] Stroud, W. J., and Agranoff, N., "Minimum-mass design of filamentary composite panels under combined loads: Design procedure based on simplified buckling equations," Tech. Rep. TN-D-8257, NASA Langley Research Center, Hampton, VA, 1976.

[91] Brown, S., "Displacement extrapolations for CFD+CSM aeroelastic analysis," *38th Structures, Structural Dynamics, and Materials Conference*, American Institute of Aeronautics and Astronautics, 1997. doi:10.2514/6.1997-1090.

[92] Secco, N. R., Kenway, G. K. W., He, P., Mader, C. A., and Martins, J. R. R. A., "Efficient Mesh Generation and Deformation for Aerodynamic Shape Optimization," *AIAA Journal*, 2020. (Submit-

ted).

[93] Luke, E., Collins, E., and Blades, E., "A Fast Mesh Deformation Method Using Explicit Interpolation," *Journal of Computational Physics*, Vol. 231, No. 2, 2012, pp. 586–601. doi:10.1016/j.jcp.2011.09.021.

[94] Sederberg, T. W., and Parry, S. R., "Free-form Deformation of Solid Geometric Models," *SIGGRAPH Comput. Graph.*, Vol. 20, No. 4, 1986, pp. 151–160. doi:10.1145/15886.15903.

[95] Kenway, G. K. W., "A Scalable, Parallel Approach for Multi-Point, High-Fidelity Aerostructural Optimization of Aircraft Configurations," Ph.D. thesis, University of Toronto, 2013.

[96] Leon, S. E., Paulino, G. H., Pereira, A., Menezes, I. F., and Lages, E. N., "A unified library of nonlinear solution schemes," *Applied Mechanics Reviews*, Vol. 64, No. 4, 2011. doi:10.1115/1.4006992.

[97] Beluni, P. X., and Chulya, A., "An improved automatic incremental algorithm for the efficient solution of nonlinear finite element equations," *Computers and Structures*, Vol. 26, No. 1-2, 1987, pp. 99–110. doi:10.1016/0045-7949(87)90240-9, URL `https://www.sciencedirect.com/science/article/pii/0045794987902409`.

[98] Belytschko, T., Liu, W. K., Moran, B., and Elkhodary, K., *Nonlinear Finite Elements for Continua and Structures*, John Wiley & Sons, Incorporated, New York, 2014. URL `http://ebookcentral.proquest.com/lib/umichigan/detail.action?docID=1501634`.

[99] Armijo, L., "Minimization of functions having lipschitz continuous first partial derivatives," *Pacific Journal of Mathematics*, Vol. 16, No. 1, 1966, pp. 1–3. doi:10.2140/pjm.1966.16.1.

[100] Grippo, L., Lampariello, F., and Lucidi, S., "Nonmonotone Line Search Technique for Newton'S Method." *SIAM Journal on Numerical Analysis*, Vol. 23, No. 4, 1986, pp. 707–716. doi:10.1137/0723046.

[101] Zhang, H., and Hager, W. W., "A nonmonotone line search technique and its application to unconstrained optimization," *SIAM Journal on Optimization*, Vol. 14, No. 4, 2004, pp. 1043–1056. doi:10.1137/S1052623403428208.

[102] Matthies, H., and Strang, G., "The solution of nonlinear finite element equations," *International Journal for Numerical Methods in Engineering*, Vol. 14, No. 11, 1979, pp. 1613–1626. doi:10.1002/nme.1620141104, URL `http://doi.wiley.com/10.1002/nme.1620141104`.

[103] Bergan, P. G., "Solution algorithms for nonlinear structural problems," *Computers and Structures*, Vol. 12, No. 4, 1980, pp. 497–509. doi:10.1016/0045-7949(80)90125-X.

[104] Sze, K. Y., Liu, X. H., and Lo, S. H., "Popular benchmark problems for geometric nonlinear analysis of shells," *Finite Elements in Analysis and Design*, Vol. 40, No. 11, 2004, pp. 1551–1569. doi:10.1016/j.finel.2003.11.001.

[105] Ghattas, O., and Li, X., "Domain decomposition methods for sensitivity analysis of a nonlinear aeroelasticity problem," *International Journal of Computational Fluid Dynamics*, Vol. 11, No. 1-2, 1998, pp. 113–130. doi:10.1080/10618569808940868.

[106] Irons, B. M., and Tuck, R. C., "A Version of the Aitken Accelerator for Computer Iteration," *International Journal for Numerical Methods in Engineering*, Vol. 1, No. 3, 1969, pp. 275–277.

[107] Sanchez, R., Albring, T., Palacios, R., Gauger, N. R., Economon, T. D., and Alonso, J. J., "Coupled adjoint-based sensitivities in large-displacement fluid-structure interaction using algorithmic differentiation," *International Journal for Numerical Methods in Engineering*, Vol. 113, No. 7, 2018, pp. 1081–1107. doi:10.1002/nme.5700.

[108] Vassberg, J., "Introduction: Drag Prediction Workshop," *Journal of Aircraft*, Vol. 45, No. 3, 2008, pp. 737–737. doi:10.2514/1.37761.

[109] Vassberg, J. C., DeHaan, M. A., Rivers, S. M., and Wahls, R. A., "Development of a Common Research Model for Applied CFD Validation Studies," 2008. doi:10.2514/6.2008-6919.

[110] Kenway, G. K. W., Kennedy, G. J., and Martins, J. R. R. A., "Aerostructural Optimization of the Common Research Model Configuration," *15th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*, Atlanta, GA, 2014. doi:10.2514/6.2014-3274, aIAA 2014-3274.

[111] Mader, C. A., Kenway, G. K., Martins, J. R. R. A., and Uranga, A., "Aerostructural Optimization of the D8 Wing with Varying Cruise Mach Numbers," *18th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*, American Institute of Aeronautics and Astronautics, 2017. doi:10.2514/6.2017-4436.

[112] Brooks, T. R., Martins, J. R. R. A., and Kennedy, G. J., "Aerostructural Trade-offs for Tow-steered Composite Wings," *Journal of Aircraft*, 2020. doi:10.2514/1.C035699.

[113] Gill, P. E., Murray, W., and Saunders, M. A., "SNOPT: An SQP Algorithm for Large-Scale Constrained Optimization," *SIAM Review*, Vol. 47, No. 1, 2005, pp. 99–131. doi:10.1137/S0036144504446096.

[114] Lyu, Z., Xu, Z., and Martins, J. R. R. A., "Benchmarking Optimization Algorithms for Wing Aerodynamic Design Optimization," *Proceedings of the 8th International Conference on Computational Fluid Dynamics*, Chengdu, Sichuan, China, 2014. ICCFD8-2014-0203.

[115] Rosenbrock, H. H., "An Automatic Method for Finding the Greatest or Least Value of a Function," *The Computer Journal*, Vol. 3, No. 3, 1960, pp. 175–184. doi:10.1093/comjnl/3.3.175.

[116] Martins, J. R. R. A., Sturdza, P., and Alonso, J. J., "The Complex-Step Derivative Approximation," *ACM Transactions on Mathematical Software*, Vol. 29, No. 3, 2003, pp. 245–262. doi:10.1145/838250.838251.

[117] Martins, J. R. R. A., and Hwang, J. T., "Multidisciplinary Design Optimization of Aircraft Configurations—Part 1: A modular coupled adjoint approach," Lecture series, Von Karman Institute for Fluid Dynamics, Rode Saint Genèse, Belgium, May 2016. ISSN0377-8312.

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

[119] Gray, J. S., Hearn, T., Moore, K., Hwang, J. T., Martins, J. R. R. A., and Ning, A., "Automatic Evaluation of Multidisciplinary Derivatives Using a Graph-Based Problem Formulation in OpenMDAO," *Proceedings of the 15th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*, Atlanta, GA, 2014. doi:10.2514/6.2014-2042.

[120] Hwang, J. T., and Martins, J. R. R. A., "A computational architecture for coupling heterogeneous numerical models and computing coupled derivatives," *ACM Transactions on Mathematical Software*, Vol. 44, No. 4, 2018, p. Article 37. doi:10.1145/3182393.

[121] Gray, J. S., Hwang, J. T., Martins, J. R. R. A., Moore, K. T., and Naylor, B. A., "OpenMDAO: An open-source framework for multidisciplinary design, analysis, and optimization," *Structural and Multidisciplinary Optimization*, Vol. 59, No. 4, 2019, pp. 1075–1104. doi:10.1007/s00158-019-02211-z.

[122] Hwang, J. T., Lee, D. Y., Cutler, J. W., and Martins, J. R. R. A., "Large-Scale Multidisciplinary Optimization of a Small Satellite's Design and Operation," *Journal of Spacecraft and Rockets*, Vol. 51, No. 5, 2014, pp. 1648–1663. doi:10.2514/1.A32751.

A

The Importance of Gradients in High-Fidelity MDO

For MDO to be useful for aerostructural design, it must be able to capture the physical phenomena responsible for constraining the design in reality, for commercial aircraft this means using computationally expensive high-fidelity RANS CFD coupled to detailed structural models. Additionally to allow for the optimiser to have control over the design at a detailed level, we want to use large numbers of design variables. These two requirements are responsible for the major challenges in the application of MDO to aerostructural design.

A.1. Optimisation Algorithms

Optimisation algorithms can be categorised by numerous characteristics, perhaps the most stark distinction is between gradient-based and gradient-free algorithms. As the name suggests, gradient-based algorithms use both the values and gradients of the objective and constraint functions at a given point to decide which point to sample next, whilst gradient-free methods require only the objective and constrain function values. Although all optimisation algorithms use some form of heuristic which pushes the solution in the direction of reducing the objective, and could therefore arguably be defined as quasi-gradient-based, it is the explicit use of gradient values which separates the two categories.

In order for it to be practical to perform MDO using expensive high-fidelity analysis, it is critical that the chosen optimisation algorithm requires the fewest possible objective and constraint function evaluations to reach an optimum, especially as problems scale to large numbers of design variables. Lyu et al. [114] compared a selection of gradient-free algorithms with gradient-based algorithms using both analytic and finite difference gradient calculation methods on the multi-dimensional Rosenbrock function, a simple but non-trivial benchmark optimisation problem [115]. Their results, shown in figure A.1, show that even for moderately sized problems with $\mathcal{O}(10-100)$ design variables, the gradient-free genetic algorithm and particle swarm methods require $\mathcal{O}\left(10^6-10^8\right)$ function evaluations, making them completely impractical for any problem where function evaluations take more than a fraction of a second to evaluate. In contrast, the gradient-based methods exhibit much better scaling with the number of design variables however, as the problem size scales, the disparity between the methods using analytic gradients and finite difference gradients grows, to extent that using finite differences is also impractical for cases with expensive function evaluations and large numbers of design variables. The only remaining choice then is to use a gradient based optimiser with an efficient method for computing analytic derivatives with respect to many design variables.

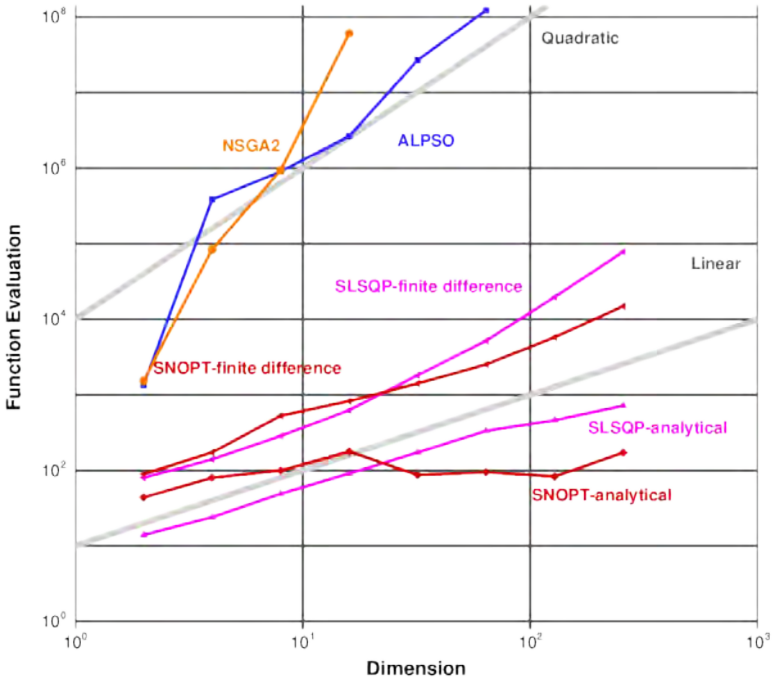

Figure A.1: Lyu et al. [114] showed that gradient based optimisation techniques using analytical derivatives require orders of magnitude fewer function evaluations than both gradient-free methods and gradient-based methods using finite differences, making them the only viable option for large scale, high-fidelity design optimisation.

A.2. Gradient Calculation

A.2.1. Perturbation Based Methods

The previously mentioned finite difference method estimates derivatives by re-evaluating a function with a small perturbation, h, to one input, using a Taylor series expansion, the derivative can be shown to be approximated by:

$$\frac{df}{dx} = \frac{f(x+h) - f(x)}{h} + \mathcal{O}(h) \tag{A.1}$$

Whilst this method requires no modification to an existing calculation code and is thus very simple to implement, it suffers from multiple accuracy problems. As shown in equation A.1, the truncation error due to the ignored nonlinear $\mathcal{O}(h)$ terms grows with step size however but, when the step size becomes too small, the derivative can be dominated by subtractive cancellation errors and possibly numerical noise in the function f. There is thus an optimal perturbation size which minimises the sum of these two errors which must be determined for a given function. This size is by no means guaranteed to be the constant across all design variables or even across the entire domain of the function, significant a priori analysis is therefore required in order to correctly tune it for a given problem. These accuracy problems can be solved by using the complex step method [116]:

$$\frac{df}{dx} = \frac{Im(f(x+ih))}{h} + \mathcal{O}\left(h^2\right) \tag{A.2}$$

Again using Taylor series, the complex step's truncation error can be shown to be $\mathcal{O}\left(h^2\right)$ superior to that of the finite-difference method and, crucially, the method does not suffer from subtractive error, making it accurate for arbitrarily small perturbations. The disadvantages of the method relative to finite-differences are that it often requires extensive code modification to allow the use of complex numbers and that it is typically significantly slower than finite-differences due to the complex arithmetic required. The bigger issue with both these methods is that, as mentioned in the previous section, both require one additional function analysis per input variable and thus render large scale MDO impractical.

A.2.2. Analytical Differentiation

The two approaches just presented are often described as 'black box' approaches as they rely solely on the inputs and outputs of a function, requiring no knowledge of the calculations performed within. In contrast, the two remaining approaches for calculating derivatives rely on implementing additional code to explicitly calculate derivatives. Analytical methods involve manually differentiating the original equations behind a piece of code and implementing new code to calculate these derivatives. This is generally the most computationally efficient and accurate method for calculating derivatives but requires a great amount of effort to implement. A particular challenge occurs when the derivative is required of a function which depends on the solution of a system of equations, such as the solution of a CFD or FE analysis, the challenge being that the calculation of such a function cannot typically be written as an explicit function. To calculate such derivatives analytically, two methods exist, the direct method and the adjoint method.

Consider a set of N_f functions of interest, $\left\{\mathcal{F}\right\}$, used for the computation of the objective and constraint functions of an optimisation problem. We want to calculate the derivatives of the function values with respect to a set of N_x design variables, $\left\{x\right\}$, also known as the Jacobian:

$$\underset{(N_f \times N_x)}{\mathcal{J}} = \underset{(N_f \times N_x)}{\left[\frac{d\mathcal{F}}{dx}\right]} = \begin{bmatrix} \frac{d\mathcal{F}_1}{dx_1} & \cdots & \frac{d\mathcal{F}_1}{dx_{N_x}} \\ \vdots & \ddots & \vdots \\ \frac{d\mathcal{F}_{N_f}}{dx_1} & \cdots & \frac{d\mathcal{F}_{N_f}}{dx_{N_x}} \end{bmatrix} \tag{A.3}$$

Each function of interest, $\mathcal{F}(x, u)$ is an explicit function of both the design variables and a set of N_u state variables, $\left\{u\right\}$. The state variables are an implicit function of the design variables through a set of residual equations which describe the physical system being simulated:

$$\underset{(N_u \times 1)}{\left\{\mathcal{R}(x, u(x))\right\}} = \left\{0\right\} \tag{A.4}$$

The Jacobian depends both on the explicit change in $\mathcal{F}$ due to the changes in x and due to the change in u required to satisfy the residual equations, using the chain rule:

$$\underset{(N_f \times N_x)}{\left[\frac{d\mathcal{F}}{dx}\right]} = \underset{(N_f \times N_x)}{\left[\frac{\partial\mathcal{F}}{\partial x}\right]} + \underset{(N_f \times N_u)(N_u \times N_x)}{\left[\frac{\partial\mathcal{F}}{\partial u}\right]\left[\frac{du}{dx}\right]} \tag{A.5}$$

Since we know that the residual equations must always be satisfied we can write:

$$\underset{(N_u \times N_x)}{\left[\frac{d\mathcal{R}}{dx}\right]} = \underset{(N_u \times N_x)}{\left[\frac{\partial\mathcal{R}}{\partial x}\right]} + \underset{(N_u \times N_u)(N_u \times N_x)}{\left[\frac{\partial\mathcal{R}}{\partial u}\right]\left[\frac{du}{dx}\right]} = \left[0\right] \tag{A.6}$$

$$\underset{(N_u \times N_u)(N_u \times N_x)}{\left[\frac{\partial\mathcal{R}}{\partial u}\right]\left[\frac{du}{dx}\right]} = -\underset{(N_u \times N_x)}{\left[\frac{\partial\mathcal{R}}{\partial x}\right]} \tag{A.7}$$

Rearranging A.7 and inserting into A.5 gives:

$$\underset{(N_f \times N_x)}{\left[\frac{d\mathcal{F}}{dx}\right]} = \underset{(N_f \times N_x)}{\left[\frac{\partial\mathcal{F}}{\partial x}\right]} - \underset{(N_f \times N_u)}{\left[\frac{\partial\mathcal{F}}{\partial u}\right]} \underset{(N_u \times N_u)}{\left[\frac{\partial\mathcal{R}}{\partial u}\right]^{-1}} \underset{(N_u \times N_x)}{\left[\frac{\partial\mathcal{R}}{\partial x}\right]} \tag{A.8}$$

The blue partial derivative terms in A.8 are relatively cheap to compute, the red term involves inverting the Jacobian of the residual equations or, in reality, solving a system of linear equations the same size as the original system of residual equations. Since the residual equations describe our high-fidelity analysis, solving

this system for a single right hand side vector can be thought of as having roughly the same computational cost as the original high-fidelity analysis. As written in A.8, calculating our Jacobian requires the solution of the system of equations for N_x right hand side vectors, in other words, one expensive computation per design variable. This is known as the direct method and clearly shows the same poor scaling with the number of design variables as the finite difference and complex step methods.

The adjoint method tackles this problem by transposing the some of the terms in A.8:

$$\left[\left[\frac{\partial \mathcal{F}}{\partial u}\right]\left[\frac{\partial \mathcal{R}}{\partial u}\right]^{-1}\right]^{\mathsf{T}} = \underset{(N_u \times N_u)}{\left[\frac{\partial \mathcal{R}}{\partial u}\right]^{-\mathsf{T}}} \underset{(N_u \times N_f)}{\left[\frac{\partial \mathcal{F}}{\partial u}\right]^{\mathsf{T}}} = \underset{(N_u \times N_f)}{\left[\psi\right]} \tag{A.9}$$

$$\underset{(N_f \times N_x)}{\left[\frac{d\mathcal{F}}{dx}\right]} = \underset{(N_f \times N_x)}{\left[\frac{\partial \mathcal{F}}{\partial x}\right]} - \underset{(N_f \times N_x)}{\left[\psi\right]^{\mathsf{T}}} \underset{(N_u \times N_x)}{\left[\frac{\partial \mathcal{R}}{\partial x}\right]} \tag{A.10}$$

A.9 clearly shows that using the adjoint requires solving the transpose of the direct methods system of equation with N_f right hand sides. Assuming an MDO problem contains a large number of design variables ($\mathcal{O}\left(10^3\right)$) and a small number of functions of interest ($\mathcal{O}\left(10\right)$), the adjoint method therefore offers an orders of magnitude reduction in computational cost over the direct method.

This derivation assumed that the governing equations of our physical system had already been discretised to form the residual equations $\mathcal{R}$, this form of the adjoint is therefore referred to as a discrete adjoint. A continuous adjoint can be formed by applying a similar technique to the underlying governing equations and then discretising the resulting continuous adjoint equations. Both methods provide their own advantages and disadvantages. The continuous method can be solved using the same solver as the original equations and is therefore typically more efficient but requires the manual differentiation of the relevant governing equations which is a significant challenge and tends to suffer from lower accuracy and other numerical issues [84]. The discrete adjoint, although less computationally efficient, has become more popular in recent years, because differentiating the discretised equations retains consistency between the behaviour of the gradient and the original function [84]. The other key advantage of the discrete approach is that the partial derivative terms (shown in blue in the above equations), do not necessarily need to be hand differentiated, instead, it is possible to compute them using automatic, or AD.

A.2.3. Algorithmic Differentiation

AD is essentially the same as the analytical approach exept that an algorithm is used to differentiate a piece of code, line by line, calculating total derivatives using the chain rule. By moving through the derivative chain from either direction, AD can be performed in forward or reverse modes, which are analagous to the direct and adjoint methods respectively. AD can be implemented using source code transformation, where the algorithm generates new code for calculating the derivatives, or through operator overloading, where elementary operations are altered to work with a new variable type which adds derivative information to a typical floating point variable. Source code transformation is generally the preffered method as it can be as efficient as analytic methods in terms of speed and memory but such tools are not available for some popular languages such as C++ [84, 117].

A.2.4. Coupled Adjoints

The power of MDO lies in its ability to correctly consider multidisciplinary tradeoffs, a task which requires the accurate calculation of multidisciplinary derivatives. Just as coupling analysis tools from different disciplines is a challenge, efficiently computing multidisciplinary derivatives represents a key obstacle to large-scale MDO. To explain this challenge, consider a system of N disciplines, each one of which is

defined by it's own system of residual equations, state variables and adjoint variables:

$$\left\{\mathcal{R}\right\} = \begin{Bmatrix} \left\{\mathcal{R}_1\right\} \\ \vdots \\ \left\{\mathcal{R}_N\right\} \end{Bmatrix}, \quad \left\{u\right\} = \begin{Bmatrix} \left\{u_1\right\} \\ \vdots \\ \left\{u_N\right\} \end{Bmatrix}, \quad \left[\psi\right] = \begin{bmatrix} \left[\psi_1\right] \\ \vdots \\ \left[\psi_N\right] \end{bmatrix}$$

Assuming that all disciplines are coupled to each other, equation A.9 then becomes:

$$\left[\frac{\partial \mathcal{R}}{\partial u}\right]^T \quad \left[\psi\right] \quad = \quad \left[\frac{\partial \mathcal{R}}{\partial u}\right]^T \tag{A.11}$$

$$\begin{bmatrix} \left[\frac{d\mathcal{R}_1}{du_1}\right] & \cdots & \left[\frac{d\mathcal{R}_1}{du_N}\right] \\ \vdots & \ddots & \vdots \\ \left[\frac{d\mathcal{R}_N}{du_1}\right] & \cdots & \left[\frac{d\mathcal{R}_N}{du_N}\right] \end{bmatrix}^T \begin{bmatrix} \left[\psi_1\right] \\ \vdots \\ \left[\psi_N\right] \end{bmatrix} = \begin{bmatrix} \left[\frac{d\mathcal{F}}{du_1}\right] \\ \vdots \\ \left[\frac{d\mathcal{F}}{du_N}\right] \end{bmatrix}^T \tag{A.12}$$

The diagonal terms of the Jacobian in A.12 are simply the Jacobians of each system of residual equations, just as appeared in the single-discipline adjoint equation. The off-diagonal terms however, describe the derivative of the residuals of one discipline with respect to the state variables of another discipline, for example the derivative of a structural solver's residuals with respect to the velocity and pressure states of a CFD code. Computing these terms adds significant complexity to the derivative calculation process as it requires the additional communication of derivatives between disciplines and the differentiation of any operations involved in coupling the disciplines. As will be described in the next section, the first coupled adjoint implementations therefore required extensive development effort and were highly specialised for the 2-3 disciplinary codes they were developed for.

Martins and Hwang [118] presented the unified derivatives equation, a common equation from which all methods for computing derivatives can be derived. Using this equation, in collaboration with Gray et al, they developed an automated method for computing the coupled derivatives of arbitrary numbers and arrangements of disciplines [119, 120] and implemented the method in NASA's open source OpenMDAO framework [121]. Hwang et al. [122] used this framework to perform a satellite design optimisation involving 7 disciplines, 25000 design variables and >2 million state variables.

B

Structural Solver Validation Cases

This appendix contains the results from 4 benchmark geometrically nonlinear shell problems used to validate TACS' nonlinear analysis capabilities. The four problems are originally defined in the NAFEMS publication 'Assembly Benchmark Tests for 3D Beams & Shell Exhibiting Geometric Non-Linear Behaviour' but have also been used and described in numerous publications since. I compare the results from TACS with the results published by Sze et al. [104] who provide results from both the commercial FE code Abaqus and, for the simpler problems, analytical theory. All analyses are performed using isotropic material properties.

Where relative errors are presented for both TACS and Abaqus, these errors are computed relative to the analytical solution data. Where only TACS errors are displayed, an analytical solution is not available and thus the errors presented are computed relative to the Abaqus solution.

In some of the problems presented, I present multiple sets of TACS results, each with a different 'drilling parameter'. By default, most shell finite elements have zero stiffness to rotations around the shell surface normal vector, often referred to as the 'drilling' degree of freedom. To avoid numerical issues, an artificial stiffness term is added to the elements which penalises any deviation of the nodal drilling rotations from the average in plane drilling rotation of the element. The scaling of this penalty term in TACS is determined by a drilling regularisation parameter. Setting this value too low can result in spurious in-plane rotational buckling modes appearing during nonlinear solutions, whilst setting it too high can degrade the accuracy of the solution by overly stiffening the structure. While performing these benchmark analyses, I occasionally found it necessary to change the drilling regularisation parameter from its default value of 10, either needing to increase it to overcome stability problems or to decrease it in order to compute accurate solutions.

B.1. Case 1: Cantilever Subjected to End Shear Force

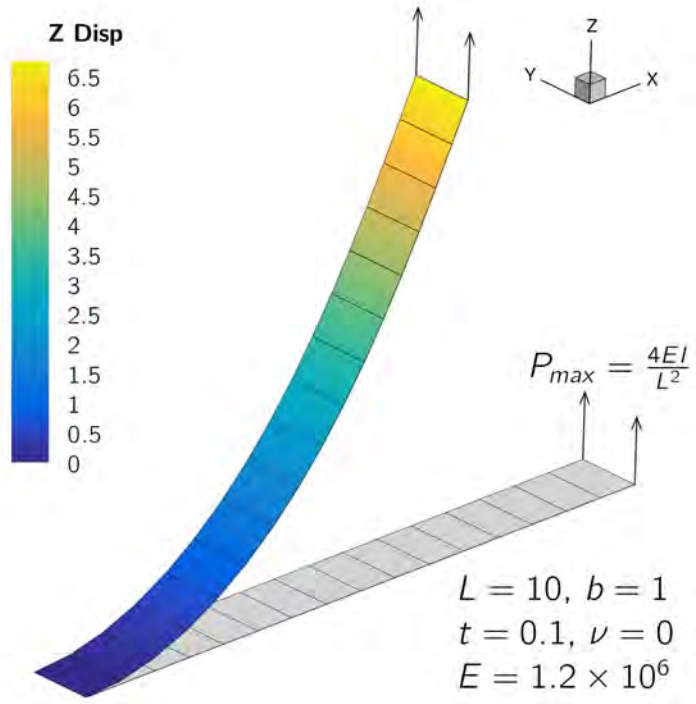

(a) Problem definition.

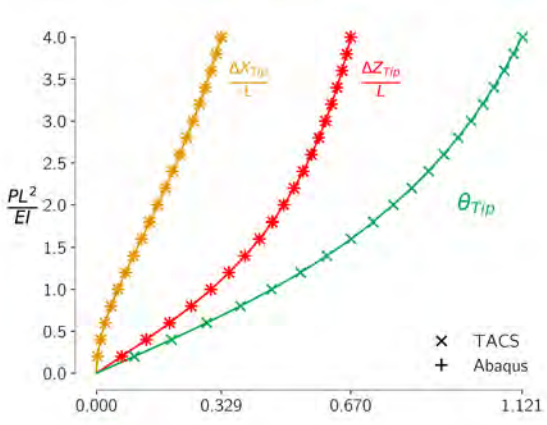

(b) Force-displacement curves.

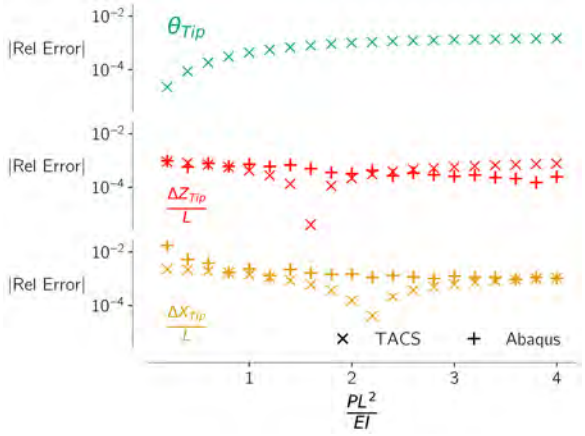

(c) Relative solution errors.

Figure B.1: Case 1: Cantilever Subjected to End Shear Force.

B.2. Case 2: Cantilever Subjected to End Moment

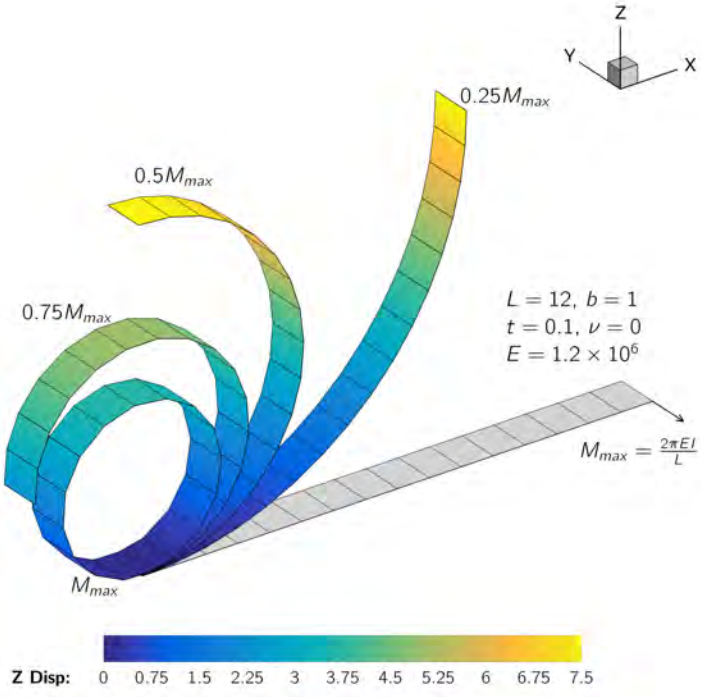

(a) Problem definition.

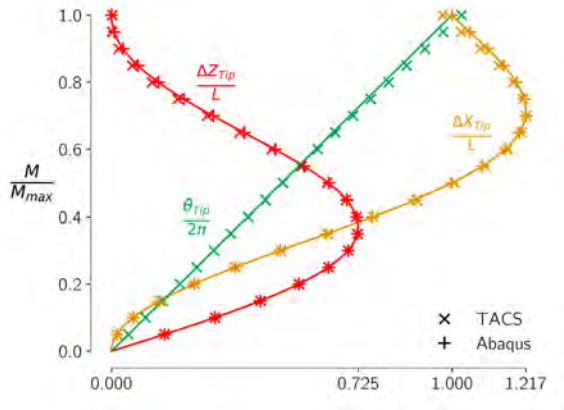

(b) Moment-displacement curves.

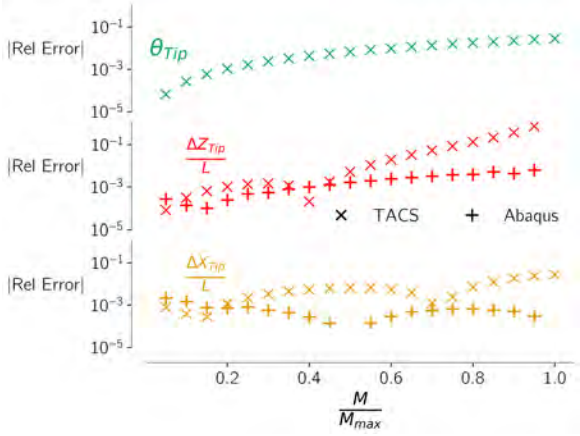

(c) Relative solution errors.

Figure B.2: Case 2: Cantilever Subjected to End Moment.

B.3. Case 3: Hemispherical Shell Subject to Alternating Radial Forces

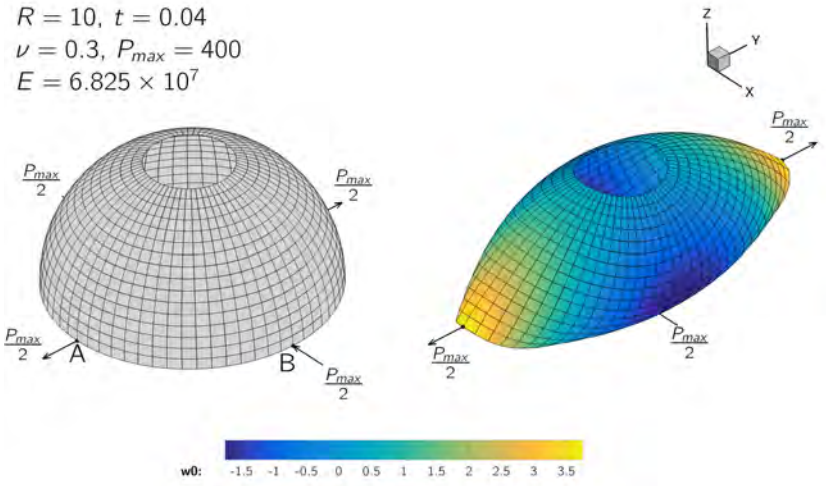

(a) Problem definition.

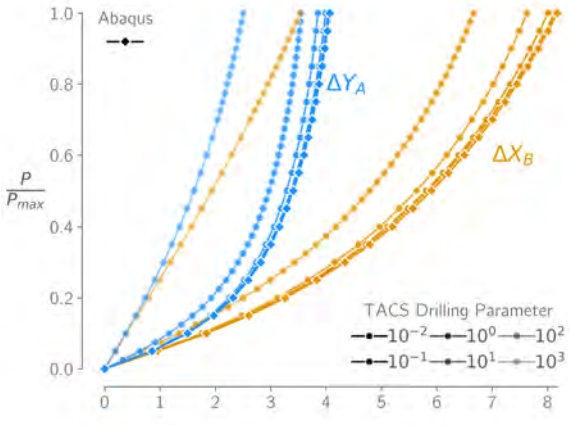

(b) Force-displacement curves.

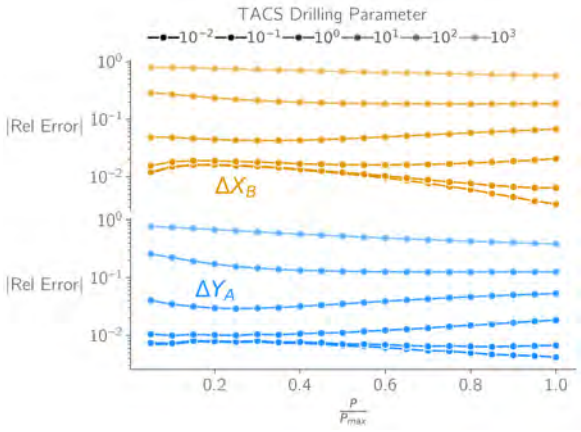

(c) Relative solution errors.

Figure B.3: Case 3: Hemispherical Shell Subject to Alternating Radial Forces

B.4. Case 4: Semi-cylindrical Shell subject to Pinching Force

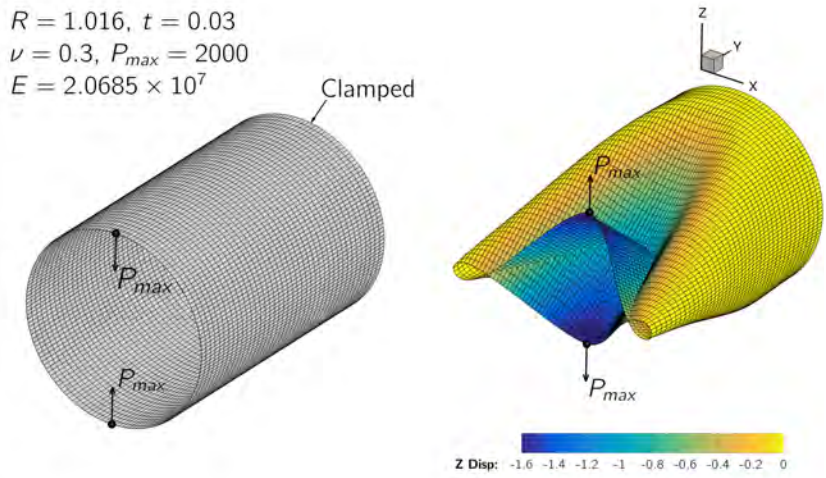

(a) Problem definition.

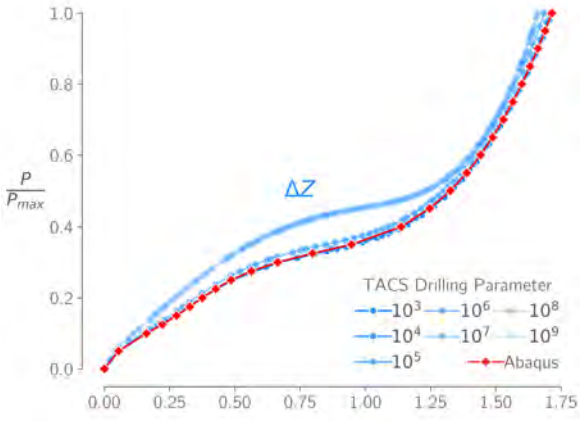

(b) Force-displacement curves.

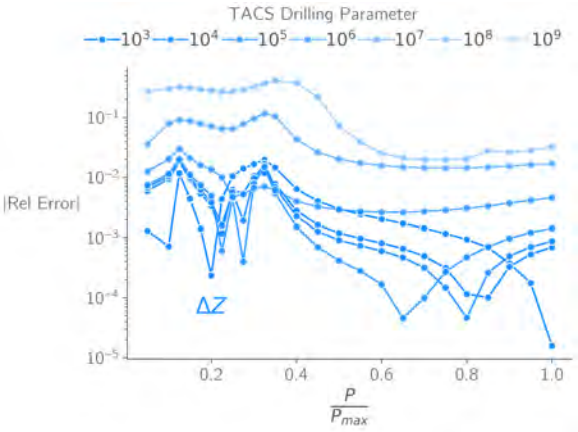

(c) Relative solution errors.

Figure B.4: Case 4: Semi-cylindrical Shell subject to Pinching Force