# Aeroelasticity of very flexible aircraft at low altitudes

**Alfonso del Carre de la Portilla**

Supervisor: Prof. Rafael Palacios

Department of Aeronautics
Imperial College London

This dissertation is submitted for the degree of
*Doctor of Philosophy*

September 2020

You could tell by the way he talked, though, that he had gone to school a long time. That was probably what was wrong with him.

John Kennedy Tool, *A Confederacy of Dunces*

# Declaration

I hereby declare that the contents of this work are original and have not been submitted, total or partially, for consideration for any other degree or qualification. This dissertation is my own work, except when explicitly stated in the text or Acknowledgements.

This work has been funded by Airbus Defence and Space and by the UK Engineering and Physical Sciences Research Council (grant number EP/R007470/1).

The copyright of this thesis rests with the author. Unless otherwise indicated, its contents are licensed under a CreativeCommons Attribution-Non Commercial 4.0 International Licence (CC BY-NC). Under this licence, you may copy and redistribute the material in any medium or format. You may also create and distribute modified versions of the work. This is on the condition that: you credit the author and do not use it, or any derivative works, for a commercial purpose. When reusing or sharing this work, ensure you make the licence terms clear to others by naming the licence and linking to the licence text. Where a work has been adapted, you should indicate that the work has been changed and describe those changes. Please seek permission from the copyright holder for uses of this work that are not included in this licence or permitted under UK Copyright Law.

Alfonso del Carre de la Portilla
September 2020

# Acknowledgements

I believe that the single most important factor in a successful PhD thesis is the people you surround yourself with and then, rather surprisingly, things fall into place (with the occasional bit of work). I had that in mind when I floated the idea of undertaking a PhD with Rafael Palacios. Now, almost five years later (four of them under his supervision) I can say that I was right: Rafa has been a brilliant supervisor, constructively steering my work through valuable advice, both professional and personal, and pushing when papers had to be written and I just wanted to do anything other than that. I hope this thesis does justice to the time and energy he devoted to me.

One of the things that attracted me about Aeroelasticity is that it lies at the intersection of many disciplines. The only way of progressing is to rely on the expertise of others in different fields. The Loads Control and Aeroelasticity Lab is proof of this. Thank you very much to everyone, present and past, who contributed to it: Rubén, Salvatore, Álvaro, Charanya, Pedro, Arturo, Yorgos and Norberto. Their influence cannot be understated and I hope my work will be as useful for them as theirs was for mine. Likewise, I would like to acknowledge the people that at some point collaborated with me. This includes the team at University of Michigan under Prof. Carlos Cesnik, as well as Dr. Keisuke Otsuka, from Tohoku University. My work is better because of them.

I would also like to deepy thank the people at the Zephyr HAPS project. It does not happen often that a company decides to fully fund three years of research for their intern. Landing an internship at the Zephyr project was a fortunate event in my career. At the time, I would have never imagined how impacted my future was going to be. I consider that a truly lucky happening. Big thanks to Alan Freeman, Michael Price and Simone Colonia.

Lastly, I would like to thank my family. My parents unconditional support cannot be understated, I hope I deserve a fraction of what they have done for me. Thanks to Paz for being my unconditional fan, and to Fiona, for dealing with me in an everyday basis, even during the tough days. Coming back at what I said at the beginning, some of the people around you are your choice, others are a result of the circumstances. But your family is given. Even if I could choose, there is not way I would have done better.

# Abstract

This work deals with the nonlinear aeroelastic response of very flexible aircraft at low altitude. In order to study it, an efficient nonlinear aeroelasticity framework, SHARPy, has been written and validated.

With it, time domain nonlinear aeroelastic simulations of a representative High Altitude, Long Endurance (HALE) very flexible aircraft immersed in 3-D low altitude boundary layer velocity fields have been conducted. Obtained loads have shown a very strong dependency with the turbulence model used in the velocity field generation. The features of the flow affecting loads the most are the span-varying character of the velocity distribution and the coherent structures that can only be captured with physics-based turbulence modelling.

Another relevant, challenging to tackle numerically, part of the low altitude segment of typical HALE missions is their reliable deployment. In order to study it, a complex simulation of a catapult-assisted take off has been set up. This highly unsteady manoeuvre is characterised by large velocity and structural deflection ranges in a short period of time. This problem requires highly robust structural dynamics capabilities, as well as an accurate unsteady aerodynamics solver.

A parametric study of the manoeuvre is included. Thanks to the improved efficiency of the code, the catapult take off problem can be optimised without the need for further simplifications. A novel hybrid optimisation algorithm based on Bayesian Optimisation and local gradient-less methods suitable for constrained expensive cost functions with low dimensionality is presented and applied to the launch problem. When optimising for a combination of ground clearance and wing root loads, the resulting manoeuvres show a strong dependency with the maximum ramp length constraint. This introduces a link between release velocity, acceleration and ramp length that affects the cost of the manoeuvre in a complex manner.

The code used in all these simulations has been released under an Open Source license.

# Table of contents

**List of figures** **15**

**List of tables** **19**

**1 Introduction** **21**

1.1 Motivation . . . 21
1.2 Aeroelasticity of very flexible aircraft . . . 23
1.2.1 High-Altitude Pseudo-Satellites: state of the art . . . 24
1.2.2 Structural dynamics of very flexible vehicles . . . 30
1.2.2.1 Linear methods for unconstrained flexible bodies . . . 30
1.2.2.2 Geometrically nonlinear beam models . . . 31
1.2.2.3 Numerical time integration . . . 33
1.2.3 Unsteady aerodynamics . . . 33
1.2.3.1 Unsteady potential flow solvers . . . 34
1.2.3.2 CFD applications in dynamic aeroelasticity . . . 37
1.2.4 Computational aeroelasticity of very flexible aircraft . . . 38
1.3 Atmospheric boundary layer modelling for linear and nonlinear aeroelasticity 40
1.3.1 Atmospheric boundary layer field generation . . . 41
1.3.1.1 Spectral turbulence generation methods for aeronautics . 41
1.3.2 Atmospheric boundary layer aeroelastic effects . . . 43
1.4 Research questions of this thesis . . . 44
1.5 Dissertation outline . . . 45
1.6 Original contributions of this work . . . 47
1.7 Publications based on this work . . . 48

**2 Numerical methods** **49**

2.1 Geometrically-Exact Composite Beam . . . 49

2.1.1 Beam kinematics . . . 50
2.1.2 Beam dynamics . . . 52
2.1.3 Multibody constraints . . . 55
2.1.3.1 Gravity loads . . . 57
2.2 Unsteady Vortex-Lattice Method . . . 58
2.2.1 Potential flow theory . . . 58
2.2.2 Discretisation and solution process . . . 59
2.2.3 Time stepping and wake shedding . . . 61
2.2.4 Force calculation from UVLM solutions . . . 63
2.2.4.1 The Joukowsky method . . . 64
2.2.4.2 The generalised Katz and Plotkin method . . . 65
2.2.5 Practical considerations . . . 66
2.2.5.1 Biot-Savart law desingularisation . . . 66
2.2.5.2 Spatial and temporal discretisation considerations . . . 68
2.2.5.3 Added mass filtering for coupled simulations . . . 70
2.3 A strongly coupled FSI algorithm . . . 72
2.3.1 Force and displacements mapping . . . 72
2.3.1.1 Force mapping . . . 72
2.3.1.2 Displacement and velocities mapping and grid generation 75
2.3.2 Relaxation factor of the FSI solutions . . . 76
2.3.3 Structural solver substepping . . . 77
2.3.4 FSI convergence check . . . 78
2.3.5 Longitudinal trim solution . . . 80

**3 Aeroelastic solver verification studies 83**
3.1 Geometrically-Exact Beam validation . . . 84
3.1.1 Nonlinear structures: dead loads . . . 84
3.1.2 Nonlinear follower loads . . . 85
3.1.3 Rigid body rotation objectivity errors . . . 86
3.2 Static aerodynamic loads . . . 87
3.2.1 3-D effects validation . . . 87
3.2.1.1 Planar wing validation against TORNADO . . . 87
3.3 Nonlinear static aeroelastic results . . . 89
3.4 X-HALE aeroelasticity . . . 90
3.4.1 The X-HALE model . . . 91

3.4.2 Structural cases . . . 94
3.4.3 Aerodynamic cases . . . 96
3.4.4 Aeroelastic cases . . . 99
3.5 X-HALE gust response . . . 101
3.6 The T-tail HALE model . . . 106

**4 Low altitude atmospheric turbulence impact on aircraft aeroelasticity 111**
4.1 Atmospheric turbulence modelling . . . 111
4.1.1 Stochastic formulations for turbulence generation . . . 113
4.1.2 High-fidelity turbulent flows via large-eddy simulation . . . 114
4.2 Comparing turbulence generation methods . . . 115
4.3 3D velocity fields as aeroelastic inputs . . . 118
4.4 Numerical results . . . 120
4.4.1 Flexible aircraft dynamics for each atmospheric turbulence model 120
4.4.2 Assessment of aeroelastic effects . . . 128

**5 A complex nonlinear manoeuvre optimisation problem 133**
5.1 Manoeuvre definition . . . 134
5.1.1 Catapult-Assisted Take Off Parametrisation . . . 134
5.1.2 Simulation process . . . 137
5.1.3 Numerical considerations . . . 139
5.1.3.1 Imposition of nodal velocity constraints for CATO simulation . . . 139
5.2 A Bayesian optimization framework for complex maneuver definition . . . 140
5.2.1 Surrogate construction . . . 141
5.2.2 Local refinement of the solution . . . 143
5.2.3 Implementation . . . 144
5.3 Optimal maneuver definition . . . 145
5.4 Numerical Results . . . 146
5.4.1 Flexibility effects on a catapult-assisted takeoff . . . 146
5.4.2 A parametric study . . . 149
5.4.3 Optimization results . . . 152

**6 Conclusions 159**
6.1 Key contributions . . . 159
6.2 Outcomes . . . 160

6.2.1 A versatile, open-source nonlinear aeroelastic solver . . . . . 161
6.2.2 Low altitude atmospheric disturbance effect on very flexible aircraft 162
6.2.3 Catapult-Assisted Take Off simulation . . . . . 163
6.2.4 Optimisation of nonlinear maneuvers . . . . . 164
6.3 Proposals for future work . . . . . 166

**References** **169**

# List of figures


1.1 Wing loading *vs* weight for birds, and solar and conventional aircraft . . 25
1.2 The NASA Helios flying wing prototype . . . . . 26
1.3 Illustration picture of Airbus-built Zephyr 8 . . . . . 27
1.4 The Solar Impulse 2 during its manned around the world flight. . . . . 27
1.5 Schematic representation of atmospheric turbulence . . . . . 43

2.1 Different frames of reference used in the description of the problem. . . . 50
2.2 Definition of the variables used to evaluate the Biot-Savart law. . . . . 60
2.3 XHALE model (see 3.4.1) solved with a free wake. . . . . 63
2.4 Biot-Savart law including the different core models implemented for this work. . . . . 69
2.5 Flowchart of the BGS strongly coupled scheme described in this work. . . 73
2.6 Force mapping procedure for coincident structural ($\boldsymbol{R}_j$) and aerodynamic ($\boldsymbol{\zeta}^{(i,j)}$) meshes. . . . . 74
2.7 Subiteration process in the strongly coupled FSI scheme. . . . . 78
2.8 Flowchart for the custom strongly coupled FSI substepping procedure. . 79

3.1 Comparison of beam shape with 2 and 10 3-noded elements for the Geradin and Cardona beam. . . . . 85
3.2 Beam shape for follower moments, including 5 and 20 3-noded elements results. . . . . 86
3.3 Comparison of obtained $C_{L\alpha}$ vs. aspect ratio for a planar wing between the current UVLM implementation and Katz & Plotkin [89] . . . . . 88
3.4 UVLM solution of planar wing . . . . . 88
3.5 Comparison of results between the current UVLM solver and Tornado for the HALE wing [175]. . . . . 89

3.6 Aeroelastically trimmed configuration with panel circulation contours for the Smith wing, $\alpha = 4°, g = 0$ . . . 90
3.7 Vertical wing deflection ($g = 0$) comparison between Smith et al. [175], and SHARPy . . . 91
3.8 Vertical wing deflection ($g$ =9.754 m s$^{-2}$) comparison between Simpson et al. [168], and SHARPy . . . 92
3.9 Photographs of the X-HALE on ground and in flight . . . 93
3.10 Nomenclature followed in the X-HALE model description. . . . 93
3.11 Spar deformation when subject to gravity. . . . 95
3.12 Spar nodal positions under a follower force ($\tilde{y} < 0$) and moment ($\tilde{y} > 0$). 97
3.13 EMX-07 airfoil used in main wing. . . . 98
3.14 Difference between the American (A) and the British (B) camber definition. 98
3.15 Spar deformation for different angles of attack . . . 100
3.16 Gust velocity time histories for different aerodynamic surfaces (as seen from the elastic axis in steady forward flight). . . . 102
3.17 $2H = 15$ chords, vertical gust response: wing spar center trajectory. . . . 103
3.18 $2H = 15$ chords, vertical gust response: body velocity in body-attached frame of reference. Long simulation (10 s in physical time). . . . 104
3.19 $2H = 15$ chords, vertical gust response: wing root pitch angle history. . . 104
3.20 $2H = 15$ chords, vertical gust response: wing tip displacement in body-attached axes. . . . 106
3.21 $2H = 15$ chords, vertical gust response: right wing root element strains. . 107
3.22 Dimensions of the HALE T-tail configuration. . . . 107
3.23 Dynamic stability diagram for rigid and flexible aircraft . . . 109

4.1 Mean velocity and turbulence intensity profiles for the three cases. The dotted lines show the reference altitude $H_{\text{ref}} = 10$ m. . . . 117
4.2 Power spectral density of velocity fluctuations for each velocity component ($u'$, $v'$ and $w'$) and all three models (1D von Kármán, 2D Kaimal and 3D LES). . . . 118
4.3 Instantaneous snapshots of a *mild* streamwise velocity, $u(t)$, for the Kaimal and LES models. . . . 119
4.4 Snapshot of a 3D LES-generated velocity field (severe case) with an overlay of a typical aircraft trajectory. . . . 120

4.5 Sample flight trajectory of the flexible configuration and all three turbulence generation methods . . . 123
4.6 Sample flight trajectory and ensemble-averaged lift/velocity distribution of the flexible configuration for the regular LES case . . . 124
4.7 Root loads from the three turbulence generation models . . . 125
4.8 Root loads envelopes for the three turbulence generation models . . . 126
4.9 Root loads envelopes for all three models obtained from combining all available time histories. In all figures the in-plane (IP) root moment is plotted against torsion. . . . 126
4.10 PDF of the incident velocity magnitude$\Delta|U_\infty|$ and out-of-plane moment $M_{OOP}$ for all models and turbulence intensity cases. Normal distribution is included for comparison. . . . 127
4.11 PSDs of the root loads (Torsion, OOP and IP) and normalised vertical acceleration at the root for all three turbulence input cases (LES, Kaimal and von Kármán) and turbulence inputs (mild, regular, severe). . . . 130
4.12 Root loads envelopes for the "Stiff" and "Flexible" cases for all levels of TI 131
4.13 Comparison of PSD distributions for the root loads (Torsion, OOP and IP) between the stiff and flexible aircraft. . . . 131

5.1 Parameters in the catapult-assisted takeoff. The aircraft goes from rest on the left to level flight on the right . . . 135
5.2 CATO parametric space including constraints . . . 135
5.3 Snapshots of the aerodynamic grid and wing tip deflections at different stages of the baseline CATO. . . . 136
5.4 Three steps of the CATO process, showing an increasing wake length as the aircraft velocity increases. . . . 138
5.5 Flowchart of the hybrid global/local optimisation process. . . . 141
5.6 Cost function contributions associated to the ground clearance (left) and wing root loads (right). . . . 146
5.7 Longitudinal trajectory for baseline CATO . . . 147
5.8 Wing root loads during baseline CATO manoeuvre . . . 148
5.9 Position of the wing tip cross section of the flexible model at different stages of the CATO . . . 148
5.10 Longitudinal trajectory sensitivity to CATO parameters . . . 150

5.11 Instantaneous cost associated to wing root loads sensitivity to CATO parameters . . . . . 151
5.12 Wing shape at the release instant under changes of release speed from baseline. . . . . 152
5.13 Contributions and total cost for a parametric study of the CATO maneuver. 153
5.14 Scatter matrix of the BO parameter sampling for three ramp lengths . . 155
5.15 Longitudinal trajectories for the optimal solutions given different ramp length constraint values . . . . . 156
5.16 Wing root loads for the optimum cases given different ramp lengths . . . 158

# List of tables


1.1 Summary of past and current HALE projects . . . . . 29

3.1 Geradin and Cardona [62] beam properties used for dead loads validation in a nonlinear context. . . . . 84
3.2 Current solver results, Simpson [168], and Geradin and Cardona [62] results comparison. . . . . 84
3.3 CRV values for a beam under a follower tip force of 3000 kN with different azimut angle between the $G$ and $A$ frames of reference. . . . . 87
3.4 Input data for the Tornado comparison case. . . . . 88
3.5 Flight conditions for the flexible HALE wing case. . . . . 90
3.6 Stiffness data for X-HALE model. . . . . 94
3.7 Distributed mass of the X-HALE model. . . . . 94
3.8 Lumped mass data. Relative positions are given with respect to the pod/wing spar intersection node. Frame of reference: $x$ right wing, $y$ upstream, $z$ up. . . . . 95
3.9 X-HALE aerodynamic model description. . . . . 96
3.10 Aerodynamic grid discretisation . . . . . 96
3.11 Structural natural frequencies (in Hz) of the aeroelastically trimmed X-HALE structural model. . . . . 97
3.12 Aerodynamic forces on the rigid aircraft, $\alpha = 4°$. . . . . 99
3.13 Flight conditions used for the X-HALE validation cases. . . . . 99
3.14 Trim results for cruise flight . . . . . 101
3.15 Stiffness and mass data for the HALE model. . . . . 107
3.16 Trim values at cruise conditions for the T-tail configuration . . . . . 108

4.1 Atmospheric turbulence flow parameters . . . . . 116

4.2 Summary of the simulations used in this chapter calculations, showing the similarity in statistics between the von Kármán, Kaimal and LES fields. . 121

5.1 Baseline for parametric study . . . . . 136
5.2 First three natural vibration frequencies of the very flexible model at different stages of the CATO. . . . . 149
5.3 Bounds for parameters of the CATO optimization problem. . . . . 153
5.4 Results of the CATO procedure optimization for different ramp length constraints and loads-associated cost scaling. . . . . 155

# Chapter 1

# Introduction

## 1.1 Motivation

The remarkable progress over the last two decades in materials science, avionics and power storage and generation solutions has enabled a new kind of vehicle: High-Altitude Long Endurance (HALE) aircraft. These configurations offer a set of capabilities [46] that were impossible to achieve a few years ago without deploying conventional satellites at a great cost. These extremely efficient unmanned air vehicles are designed mostly to achieve flight times of the order of weeks or months at a time by charging their batteries and gaining altitude during the day and being extremely energy-efficient at night. This inevitably results in slender, very light structures that can present a very particular set of challenges from an aeroelastic point of view. Nowadays, the development of HALE configurations is driven by the so-called High-Altitude Pseudo-Satellite (HAPS) application. The idea behind this is to use constellations of quasi-perpetual HALE aircraft to provide a lower cost alternative to conventional satellites. Other advantages are worth considering when discussing HAPS capabilities: quicker redeployment, possibility of payload reconfiguration, and simpler payloads due to the lower altitude at which they operate, with less stringent environmental conditions.

In a given HAPS mission, the time spent at high altitude is orders of magnitude larger than the time subject to the lower atmospheric boundary layer. However, the reduced number of HALE-type aircraft flights conducted until now have shown that the aircraft response to lower altitude atmospheric disturbances can be much more important than it would be expected if only relative mission time is considered. The typical HAPS

configuration is a low speed, lightweight, span-loaded aircraft with reduced wing loading, high aspect-ratio wings and overall low controllability.

Assessing the potentially complex response of these non conventional aircraft is challenging from several points of view. The lightweight structures of these aircraft result in low structural natural frequencies. When these frequencies are low enough, they can overlap with the rigid body modes frequencies, which is most likely the case for all modern HALE-like aircraft. This potentially introduces complex interactions between the structural and the flight dynamics which are not generally captured with conventional aeroelastic methods. Other effects that can be inaccurately predicted with industry-standard methods are the changes in mass and stiffness matrices due to nonlinear deflections and the change in aerodynamic shape, both being potentially important in the overall analysis of the platform. An interesting effect that illustrates how these flexible, large span aircraft can show different behaviour than predicted in conventional aircraft is the load factor/wing bending relation. In certification bases for small [53] and large [52] aircraft, structural loads are assumed to be a linear dependency of the load factor. Large span, flexible aircraft show a different trend, where a gust encounter at high speed can produce large $g$ loads, while the bending moment of the wing root is smaller than at low speeds. This is due to the reduced torsional stiffness of the long, slender wing spars, together with the typical HALE airfoils, where the chordwise pressure centre moves aft with increasing speeds.

Industry-standard methods are mostly based on frequency domain analysis, linearised models with partial or total decoupling of the rigid body motions and the flexible structural dynamics and some higher-order corrections to the models. While their suitability for conventional, faster, heavier and stiffer aircraft is beyond any doubt, the application of these methods for typical HALE configurations cannot be taken for granted.

A typical solar-powered HALE configuration has a large fraction of the total mass distributed along the wing span in the form of battery packs. This results in load distributions under longitudinal gusts that greatly vary with respect to conventional aircraft, where a large portion of the mass is in the fuselage or close to the wing root. Moreover, the centre of mass variation in position when wings undergo nonlinear deformations can be important, due to a large ratio of the mass being attached to the wings. Finally, spanloaded lightweight airframes enhance the role of inertia relief as a passive load alleviation mechanism. To date, most of the computational effort in the open literature has focused on building nonlinear aeroelastic simulation strategies to capture those effects [138, 202, 214, 180, 70] and, more recently, on designing suitable feedback

control strategies [66, 29, 163, 143, 9, 200, 211]. An additional challenge is brought about by the need to identify nominal operation procedures, as recently explored by Maraniello and Palacios [107] for lateral maneuvers with very flexible wings. A parametrisation of the maneuver definition defines an optimization problem on a suitable metric of the dynamic response of the vehicle.

This dissertation is an investigation into the numerical methods and aeroelastic analysis tools necessary for the development of a next generation of these platforms. The included methods and implementation details, as well as the problems presented in this work are the result of a long term effort aiming to identify and solve potentially critical gaps in knowledge and simulation capabilities. The use of the resulting simulation tools in an optimisation framework adapted to the particular properties of this kind of problems is explored as well. This work aims to accurately describe the current challenges arising from nonlinear aeroelastic simulations of HALE aircraft, while providing a reasoned argument about the possible solutions with their strengths and weaknesses.

## 1.2 Aeroelasticity of very flexible aircraft

Aeroelasticity has been an important topic on aircraft design almost since the first prototypes. Already in 1906 there was an interest in static divergence phenomena, as lightweight structures and lack of understanding resulted in torsional divergence problems (Samuel Langley's second prototype failed due to it). Tail flutter was first documented in 1916 (in the Handley Page O/400 bomber). However, the term aeroelasticity was not coined until the early 1930s by scientists at the Royal Aircraft Establishment (RAE) in Farnborough, UK. Today, our understanding of the subject has~~, thankfully,~~ greatly improved (see Refs. [207, 12, 57]). Modern aircraft are designed with a multitude of requirements derived of the lessons accumulated this last century: from freeplay and backslash in control surface actuators to the stiffness of every structural element of the aircraft are designed with these requirements in mind.

Newer commercial aircraft such as the Airbus 350 or the Boeing 787 feature relatively flexible wings, with relative wing tip displacements close to 10% of the semispan during cruise flight. Typical HALE platforms are the next step in terms of flexibility, with wing tip displacements even larger, sometimes reaching 35%. These will be identified as Very Flexible Aircraft (VFA) during this work. VFA aeroelasticity requires nonlinear analysis [138], which implies a need for different tools and modelling approaches in order to accurately capture the phenomena relevant to this type of aircraft. In a certain way,

the design philosophy of HALE aircraft is closer to the early prototypes being developed in the early 20th century: very light construction, relatively flexible structures, low controllability and low flight speeds. Obviously, the main difference between these is the intention behind the design. While early prototypes were about experimenting and determining what worked and what did not, recent HALE aircraft make use of new advances in many different disciplines in order to maximise endurance and push the envelope of what is considered possible nowadays. This also involves a great deal of experimentation and out-of-the-box thinking, as the particular flight regime and properties of HALE configurations pose a set of challenges that might not overlap with the common ones in conventional aircraft aeroelasticity. Figure 1.1 shows a wing loading/total weight diagram for solar aircraft amongst other kinds of flying animals and conventional aircraft at the time (2006). It constitutes a good visualisation of how particular HALE platforms are. Since then, new developments have pushed even further the HALE platforms towards the upper left corner of the envelope, with higher weights and lower span loading, which unavoidably results in larger lifting surfaces.

Numerical methods for very flexible aircraft aeroelasticity need to account for large deformations, structural cross-couplings and rigid body motion interaction with the elastic degrees of freedom. All of these phenomena need to be integrated in a framework where aerodynamic loads capture the nonlinear structural deformations and the coupled character of their mutual interactions. In short, a valid nonlinear aeroelastic framework needs to integrate a robust structural model of free very flexible bodies, an aerodynamic solver that supports large and possibly fast deformations and a coupling mechanism that accurately captures their interaction. This literature review provides an overview of the last two decades in HALE aircraft engineering, as well as the history and recent developments in these three fields.

### 1.2.1 High-Altitude Pseudo-Satellites: state of the art

Since the turn of the millennium, a relatively large number of companies have ventured into the development of these solutions, with different levels of success. This section gives an overview of the state of the art in early 2020, starting with the better known HAPS project in the aeronautical community: the NASA Helios. This is followed by an overview of active projects, and the section finishes with a few cancelled projects. A more in-depth work about this, although slightly outdated, can be found in Ref. [46].

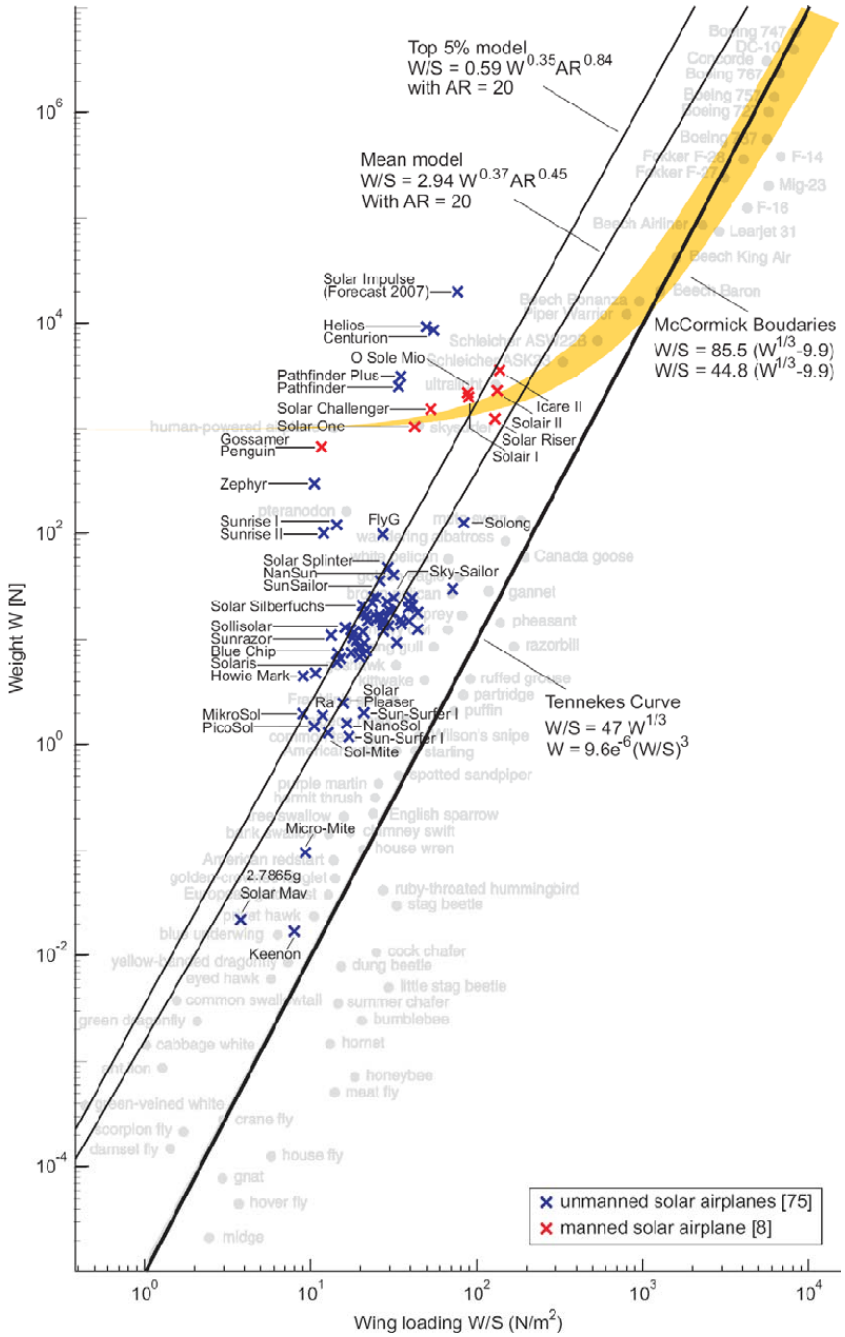



**Fig. 1.1** Wing loading *vs* weight for all kinds of flying bodies. The solar aircraft included in this plot in blue and red crosses are from a decade ago. Current configurations have an even lower weight. Extracted from [190].




**Fig. 1.2** The NASA Helios flying wing prototype on a test flight on July 14, 2001. Photo by NASA

The first serious efforts to develop a full-sized HALE were made by NASA and AeroVironment. The result of this work was the Pathfinder and later Helios (Fig. 1.2) platforms, a relatively span-loaded very flexible configuration with multiple pods attached to a straight single wing. Initially, the span-loaded version of the aircraft, the HP03-1, had been analysed and tested successfully. Later requirements resulted in the addition of extra mass in the central pod —the HP03-2—, this resulting in a different aeroelastically trimmed state. While the aeroelastic analysis tools used at the time did not show any unstable behaviour, the atmospheric turbulence resulted in an increased dihedral trim point, which modified the stability derivatives of the platform, causing the aircraft to enter an undamped phugoid motion, where the airspeed oscillated while steadily increasing until maximum airspeed was reached and surpassed. This resulted in catastrophic structural failure. The report of the incident [130] identified the root causes of the accident. The most relevant for this work is the lack of (at the time) adequate multidisciplinary analysis tools that could account for the potentially large deflections occurring in-flight.

The company behind the Helios design (AeroVironment) has recently joined SoftBank to create HAPSMobile, with the target of developing the HAWK30, a development of the older Helios. The HAWK30 is a tailless flying wing weighting around 1000 kg (the HP03-1 had a payload of 329 kg) with a 78 m span and 10 propellers. This aircraft has been flown at low altitude and a high-altitude flight at 65000 ft is planned for 2020.

Overlapping for a few years with the Pathfinder/Helios, QinetiQ started developing the Zephyr platform in 2001 [146]. Later in 2013 the Zephyr Project was purchased by Airbus. The last operational version of Zephyr, called Zephyr S (S for single tail), is the endurance flight record holder with an almost 26 days flight. Zephyr S weights less than 75 kg in total, with a span of 25 m. Another version of Zephyr, Zephyr T (for twin-tail)



**Fig. 1.3** Illustration picture of Airbus-built High Altitude Pseudo-Satellite Zephyr 8. Image by Airbus Defence and Space.



**Fig. 1.4** The Solar Impulse 2 during its manned around the world flight.

will feature a greater span and total weight. Zephyr S is also the first HAPS service to be provided in a commercial contract.

In September 2019 BAE Systems acquired Prismatic, the company developing the PHASA-35. It is a single tail, twin propeller configuration with a span of 35 m, a total weight of 150 kg and a payload of 15 kg. A Reynolds number-scaled vehicle (PHASE-8), with a 8.75 kg total mass has been flown in 2017 and flight tests for the PHASA-35 are scheduled for early 2020.

In China, CAAA has developed the Caihong, a twin-tailed aircraft with 8 propellers and a span of 45 m. The prototype has already been flown for one day at 65000 ft.

The Solar Impulse 2 completed in 2016 the first manned circumnavigation of the globe on solar energy. It also is the holder of a number of records related to endurance and range of solar aircraft. The newest active HAPS to start operations has been Skydweller Aero. They acquired the Solar Impulse 2 (Fig. 1.4) and it will be used as the base for a heavy payload HAPS solution. The Solar Impulse 2 weights 2400 kg, with a span of 72 m. They intend to fly an Optionally Piloted Vehicle (OPV) in 2020 and the first fully unmanned flight is scheduled for 2021.

Amongst the prototypes still in development, the Russian Lavochkin Laboratory is developing the LA-252: a 25 m span, 116 kg HAPS solution designed for a mission length of 100 days, while Ordnance Survey in the UK, in partnership with Astigan is developing

the A3, a 38 m span, 149 kg total mass (25 kg payload) HAPS with an endurance of 90 days.

Most of the previous projects are relatively recent, being created in the last decade. During that same time, others started and were indefinitely delayed or cancelled. The Solar Eagle, also called Vulture II, was developed by Boeing Phantom Works coordinated by DARPA. The now cancelled project had a total mass of 2700 kg and a span of 121 m. Facebook purchased the British company Ascenta. Their Aquila model was a flying wing with a span of 42 m and a total weight of 400 kg. A 1:10 scale model was flown in 2014, but the project was later cancelled after logistics-related problems and a crash resulting to structural failure after a low-altitude gust hit the starboard wing. Aurora Flight Sciences also started to develop a HAPS solution. The Odysseus was planned to have a span of 75 m and a payload mass of 25 kg. The geometry is based on a straight wing with three tails (the outer ones providing roll control) and six propellers. The project got indefinitely delayed in July 2019.

Table 1.1 provides a summary of the past and current HALE projects together with relevant information about payload and dimensions.

An interesting point to note is that the four companies that developed a full-size HALE aircraft prototype (NASA/AeroVironment, QinetiQ/Airbus, Google and Facebook) have suffered incidents that resulted in structural damage due to unexpected turbulence loads. While the long-endurance/high-altitude segment of the mission is the one currently driving the slender, lightweight designs, it is clear that the effect of low-altitude atmospheric disturbances on the structural loading cannot be dismissed and structural, flight dynamics and control system responses under unexpected turbulence need to be considered from the very early stages of design in a holistic manner.

| **Company** | **Project** | **Years** | **Payload/ Total weight [kg]** | **Span [m]** | **Status** | **Notes** |
|---|---|---|---|---|---|---|
| NASA/ AeroVironment | Helios | 1999–2003 | 329/1000 | 75 | Cancelled | Crashed in 2003 |
| AeroVironment | HAWK30 | 2018– | | 78 | Active | First flight scheduled for March 2020 |
| Airbus Defence and Space | Zephyr | 2003– | S - 5/75<br>T - /140 | S - 25<br>T - 33 | Active | World record holder: almost 26 days<br>First commercial contract for HAPS services |
| BAE Systems/ Prismatic | PHASA-35 | 2017– | 15/150 | 35 | Active | Quarter-size model flown in 2018 |
| CAAA | Caihong CH4 | – | /400 | 45 | Active | Flown at 66000 ft for 1 day |
| Skydweller Aero | Solar Impulse 2 | 2019– | 360/2300 | 71 | Active | Solar Impulse 2 completed a manned tour around the world |
| Lavochkin Laboratory | LA-252 | – | /116 | 25 | | Active on January 2018 |
| Ordnance Survey/ Astigan | A3 | 2014– | 25/149 | 38 | Active | |
| Boeing Phantom Works/ DARPA | Solar Eagle/ Vulture II | 2010–2012 | /2700 | 121 | Cancelled | |
| Ascenta/Facebook | Aquila | 2014–2018 | /400 | 42 | Cancelled | Crashed (structural failure) during first flight |
| Aurora Flight Sciences | Odysseus | –2019 | 25/ | 74 | Indefinitely delayed | |

**Table 1.1** Summary of past and current HALE projects

### 1.2.2 Structural dynamics of very flexible vehicles

Development of lightweight spacecraft encouraged the research into Flexible Multibody Dynamics (FMD) in the second half of the past century. Flexible Multibody Structures are structures subject to potentially large translational and rotational deformations, in a constrained or unconstrained environment.

FMD analysis is typically treated as a composition of two different motions. First, the multibody analysis of the body as a rigid body describes its position with respect to a certain inertial frame of reference. The second part of the motion is the elastic deformation contribution, which can be described in a number of different frames of reference. From Shabana [158, 159, 157] and Wasfy and Noor [204, 205], the types of frame of reference used to described the elastic contribution can be classified in three categories: floating, co-rotational and inertial. The floating frame approach consists on splitting the motion of structures that experience deformations into a frame describing the mean rigid body motion dynamics, and a superposed motion that defines the flexible deformation state. It is suitable for small —linear— deformations, while its application to nonlinear deformations does not provide any benefit [112]. The co-rotational approach consists on separating rigid body motions from strain-producing deformations at local level. To accomplish this, a local frame of reference is attached to a node of the structure, which then rotates and translates with it. This formulation supports geometrically nonlinear deformations, as long as the material constitutive equations are linear. The inertial description describes the motion of the elements without an intermediate frame of reference for the local deformation.

Evidence from the literature shows that for geometrically nonlinear elasticity, it is convenient to use the inertial approach. This allows to define a reference for the nodal orientations.

#### 1.2.2.1 Linear methods for unconstrained flexible bodies

A first approximation to the flexibility effects on the response of the aircraft can be obtained with linear methods. These methods define a floating frame of reference in order to superimpose the elastic deformations over the rigid body motion of the body. While some authors [158, 21] have investigated the use of body-attached frames of reference in a linear context, the fact that the moving axis formulation allows to minimise the magnitude of the small (linear) elastic deformation makes them a very attractive option for linear analysis [18, 19, 10] by keeping linearisation error small.

The common choice for a moving frame of reference is the Tisserand frame [195, 18, 131], which effectively minimises the kinetic energy of the elastic motion of the moving frame. As a consequence of this, the Tisserand frame of reference fulfils the mean axis approximation, meaning that the floating frame of reference balances the momenta due to the elastic deformations. The author that first described this for aircraft applications is Milne [117, 118], who shows that under small enough displacements, there exists a time-dependent frame of reference which decouples the flexible and the rigid-body degrees of freedom. The resulting formulation is greatly simplified due to this, and has been applied in spacecraft [4, 18, 23], and aircraft dynamics [56, 93, 117, 206].

An important consideration is that the mean axes approximation is enforced when using free-free modes in an unconstrained structure. This is due to the position of the centre mass being preserved when subject to small amplitude oscillations. Furthermore, the modal projection of the flexible-body dynamics allows to truncate the higher frequencies modes, reducing the complexity of the model and allowing the use of experimental model identification methods [33, 3].

A linear undamped system has orthogonal normal modes. Thus, it can be reduced to a set of ODEs. However, when arbitrary motion, specifically large angular rotations, is introduced, we need to account for the Coriolis and centrifugal forces. The coupling between the rigid-body and flexible motion that these forces introduce is neglected when mean axes approximations are used, as Meirovitch proved in Refs. [111, 112].

#### 1.2.2.2 Geometrically nonlinear beam models

Better composite material understanding and expertise has encouraged its utilisation for load-bearing structures. This has resulted in reduced weight, more flexible structures, which in turn, have pushed forwards the development of geometrically nonlinear composite beam models. They consist on two steps: an homogeneisation, where the three-dimensional composite structure is reduced to a one-dimensional beam model [26, 45, 90], and the solution of the 1-D dynamic equations of motion of the homogenised beam [10, 73, 74]. Hodges provides in [75] a comprehensive review of the literature on composite beam modelling.

Geometrically nonlinear beam theories are commonly based on the inertial frame approach, where the beam dynamics are described with respect to a reference frame that incorporates both structural deformations and rigid-body motion in one set of equations. The geometrically-exact theory requires that the cross-sectional orientation is represented by a rotation parametrisation. Several families of parametrisation are available. Euler

angles were used by Reissner [149], while Simo and Vu-Quoc [165, 166] generalised the approach to the fully dynamic 3-D case by using quaternions to describe the rotational degrees of freedom. Later, these same authors extended their theory to rods under large unconstrained motions in space, including rotations [167]. A global body-fixed frame of reference attached to the reference configuration was used to include the rigid-body motion of the beam. This geometrically-exact description was based on nonlinear strain measures, accurately capturing axial, shear, torsion and bending deformations.

Géradin and Cardona used the Cartesian Rotation Vector parametrisation for the section orientation of the beam elements [20, 62]. Doing so is advantageous for a number of reasons: first, the CRV as a representation of the rotation tensor requires less storage than the quaternion parametrisation used by Simo and Vo-Quoc. More importantly, the CRV can be easily linearised using derived expressions. Further work was devoted to improve the rotations parametrisation in order to extend the formulation to initially curved beams [80].

All the nonlinear formulations presented up to this point have used nodal displacements and rotations as their primary variables. However, more recent work [24, 161] has focused on strain-based formulations to solve the geometrically-exact beam equations using the decomposition of the beam strains presented in [35]. This formulation can be then rearranged into state-space form with a constant stiffness matrix, which increases numerical efficiency. However, in order to obtain displacements and rotations, a post-processing stage is required [133]. Hodges [73] developed a hybrid formulation, called intrinsic, where beam strains and local translational and angular velocities are used as the degrees of freedom of the problem. Ref. [133] provides a comparison between structural and aerodynamic models for the modelling of flexible aircraft dynamics. In that work, the strain-based and intrinsic formulations efficiency advantages were confirmed, while the ease of coupling with the aerodynamic model and boundary condition implementation were highlighted as advantages of the displacements and rotations formulations.

Hodges' intrinsic formulation guarantees rotation objectivity, which is the invariance of the strain measure with respect to arbitrary (and possibly very large) rigid-body motions [73]. It can be proven that displacement-based beam formulations cannot be objective in the interpolation of the rotation degrees of freedom because of the inertial description of the beam kinematics, as Ref. [34] shows. Despite this advantage, this work proposes to apply instead the inertial beam formulation. This would allow to capture possible large beam deflections with the application of simplified strain measures [78]. On the other hand, Ref. [168] demonstrates that the error in the strain field arising from

interpolation rotations decreases with $p$- or $h$-refinement, ensuring that displacement-based formulations can be in practice applied in the dynamics of very flexible structures. In practice, it has been shown [168] that for nonlinear aeroelasticity with coincident discretisation, the aerodynamic grid is the main driver for mesh refinement. Thus, the numerical error introduced by the lack of objectivity of the rotation interpolation stays under the error of the method.

#### 1.2.2.3 Numerical time integration

The second-order ordinary differential equations set that results from the spatial discretisation of the geometrically-exact composite beam equations is in most cases numerically stiff. This is due to the presence of rigid-body motion degrees of freedom and the broad frequency range of the response. In an effort to improve the numerical stability of the structural solver, implicit numerical time integration schemes with unconditional stability properties are chosen. The original Newmark method [129] is a common starting point and the most used time integration algorithms in the FMD community due its ability to damp the high-frequency numerical artifacts that arise in these solutions. The generalisation to nonlinear problems using Newton-Raphson implicit subiterations was proposed by Géradin and Rixen [63]. A similar method, by Shearer and Cesnik [160], has been used in this work.

Energy-preserving schemes have been extensively studied in the literature [99, 95] due to their suitable intrinsic properties. Ref. [96] studies the energy-preserving character of the explicit Newmark. The necessary addition of numerical damping in the higher-frequency modes results in the loss of its second-order accuracy. Other methods such as the Hilber-$\alpha$ [71] achieved second-order accuracy with unconditional stability, as well as higher-frequency damping, at the cost of more complex numerical procedures, which results in an added overhead if compared to the Newmark-$\beta$ method. While useful, its second order accuracy is not a requirement for coupled simulations, as the time step of the simulation is usually restricted by the aerodynamic solver and not the structure, as long as numerical damping of the higher structural frequencies is possible.

### 1.2.3 Unsteady aerodynamics

The available computational methods for solving unsteady aerodynamics are varied in complexity, accuracy and range of validity. The classical fundamental tool for aeroelasticians has been the Doublet-Lattice Method (DLM) [6], which is a frequency-domain

solver, involving some linearisations. Other methods based on strip theory are still popular for preliminary design stages [92], where the cost has to be low enough to allow for a large number of evaluations. On the other hand, high-fidelity Computational Fluid Dynamics (CFD) methods, which are based on a discrete approximation of the (usually simplified) Navier-Stokes equations, such as the Unsteady Reynolds-Averaged Navier-Stokes (URANS), are being increasingly used in industrial environments, specially when dealing with transonic flow [110, 64, 37]. For HALE applications, URANS solvers can be useful for more accurate drag prediction and overall performance and stability when approaching the stall regime.

Between these two extrema, other methods are available, although obviously, the higher the accuracy, the higher the computational cost. Due to extensive experience with potential flow solvers, the industrial loads analysis for conventional aircraft still relies mostly on lower-fidelity aerodynamic solvers. Furthermore, low or medium-fidelity aerodynamics are the accepted certification procedure for gust load cases, as the CS-25 [48] Certification Specification AMC (Acceptable Means of Compliance) states.

#### 1.2.3.1 Unsteady potential flow solvers

HALE platforms usually operate on a regime characterised by attached, sometimes laminar, flow at low-speed, with a Reynolds number around $10^5$, which is relatively small if compared to conventional aircraft, but high enough to obtain pressure distributions that are accurately described by inviscid flow models [27]. Potential flow solvers are specially well suited to solve this kind of flow dynamics, while other higher-fidelity methods provide more accurate drag estimation and include stall effects. A brief review of the available potential flow solvers is included below.

Incompressible potential flow solvers incorporate the time dependence through the boundary conditions [89]. Although the fundamental equations are linear, the nonlinearities can appear from large changes in geometry. These geometric nonlinearities are a consequence of a displacement of the body panels that cannot be simplified into a rigid-body motion, e.g. a wing deformation, or the influence of the free wake over the induced velocities on the lifting surfaces [203]. These deformations are accurately captured as long as the flow around the surfaces remains attached.

A common classification for potential methods is based on the kind of singularity used. 2-D and 3-D solvers normally use sources, vortices, doublets or a combination of them. It is important to note that higher-order panel methods exist, where the distribution of singularities is not constant throughout the associated panel. With these methods, fewer

panels would be necessary to obtain a certain precision with a higher-order formulation. In practice, the geometry cannot generally be defined accurately with large panels, and more panels are needed anyway. The community agrees that low-order methods are a good combination of cheap and effective [108, 89], while very few high-order implementations have proven to be worth it [197].

3-D solvers are the main focus of this discussion due to the importance of predicting the potentially time-varying spanwise lift distribution correctly, which can only be obtained with 3-D models, even in the case of high aspect-ratio wings. The spanwise life distribution due to 3-D effects has a very strong influence in the steady loads and the unsteady dynamics of the aircraft, especially once flexible structures are considered.

Unsteady lifting-line theory [5, 47] is a modification to the widely-used lifting-line method. It yields a first approximation to the spanwise variations in a potential flow solver due to finite wing effects. The unsteady modification accounts for the wake convection with two-dimensional empirical lag terms [47]. Unsteady lifting-line theory captures finite wing effects, large deformations and surface-surface interactions. However, it is only suitable in the very high aspect-ratio, slowly moving surfaces case. Apart from flutter speed validation for extremely aspect-ratio wings ($\mathcal{AR} \approx 200$, basically a typical section problem), there has been, to the knowledge of the author, no published work that validates flutter predictions for realistic 3-D wings using this model.

Nowadays, lifting surface methods are the most widely used in aircraft aeroelastic analysis. As stated earlier, the industry standard for conventional aircraft aeroelasticity is the Doublet-Lattice Method (DLM), although other alternatives exist, such as the Unsteady Vortex-Lattice Method (UVLM) [89]. Lifting surfaces are based on placing arrays of polygonal —usually quadrilateral— elements over the mean lifting surface of a model. The DLM is a frequency-domain, very effective method for relatively stiff aircraft. The non-penetration condition is enforced over the undeformed geometry [13]. It is due to this that the DLM is not suitable for flexible aircraft analysis where, even in steady flight, large deformations are common. From these boundary conditions, the circulations of the potential singularities attached to the panels are calculated.

DLM can be used for obtaining time-domain models. Rational Function Approximations (RFA) are usually obtained using Roger's method [151, 16, 41], which is based on Padé approximants. Only an inverse Laplace transform is then necessary to obtain a time-domain model.

Finally, the Vortex-Lattice Method (VLM), is effectively a DLM at reduced frequency, $k = 0$ [150], as long as the incidence angle of the surface is small, due to linear approxi-

mations embedded in the DLM formulation [6]. The common ground between solvers disappears as we start considering unsteady analysis, when the reduced frequency $k = \frac{\omega c}{2V}$ becomes larger than zero.

**The Unsteady Vortex-Lattice Method** The UVLM is a 3-D, lifting surface method which deals with arbitrary kinematics of the lifting surfaces. Aerodynamic bodies are represented by their mean aerodynamic surfaces (an example can be seen in Fig. 2.1). Each panel from the discretisation has a vortex ring associated with it. A complete recent description of the method is given in [89]. Later works have expanded with discussions about induced drag calculation [168], numerical implementation [28] and acceleration methods [192, 14]. The standard UVLM does not consider non-lifting bodies, and it is limited to thin wings in an attached flow condition.

The UVLM yields a time-domain, naturally discrete system of equations. The non-penetration boundary condition is optionally enforced on the deformed configuration at every step, in contrast with the DLM, where the AIC (Aerodynamic Influence Coefficients) matrix is kept fixed. This results in a geometrically-exact representation. Another important difference between these two methods is that the DLM assumes a flat wake even in the unsteady case, where the UVLM supports wake convection, and even free wake simulation following Helmholtz's laws of vortex motion.

Free wake nonlinearities can be important in unsteady simulations even in simple geometries. Suciu and Morino [181] shows that flat wake or free wake modelling in steady simulations only presents marginal differences between them in both spanwise lift distribution and overall lift. However, neither the effect on the drag or the free-wake nonlinearity in an unsteady motion was considered. Wake interference effects in HALE configurations have also been studied by Murua et al[124]. Results showed that a prescribed wake induces sinusoidal variations in tail lift when simulating a phugoid motion. The authors state that it was not possible to obtain very detailed results, as UVLM became too expensive for a long free wake for that particular implementation. Independently of this, the coarse discretisation used in that work proved that wake roll-up can have a strong effect on the tail lift and this, special attention has be to payed to the wake modelling when simulating full configurations (wing and tail) in unsteady situations.

#### 1.2.3.2 CFD applications in dynamic aeroelasticity

For a long time, steady and unsteady RANS simulations of flexible aircraft were sparsely used due to their cost. Even nowadays, RANS is still too expensive to routinely use in the aeroelastic design or certification stages, where hundreds or thousands of cases need to be run.

The workaround to the limited computational resources available compared to the cost of a full URANS simulation was first to increase the fidelity of the aerodynamic model from a panel-based solver to Euler methods. This offered an improvement in transonic conditions prediction (this work will not consider any compressibility effects due to the characteristic flight regime of HALE aircraft). Nowadays, the most common way of performing higher-fidelity aeroelastic analysis is to correct the aerodynamic influence matrices —usually obtained with a DLM [144]— with a higher-fidelity method such as the Linearised Frequency Domain (LFD) method [187]. LFD methods retain the accuracy of the full RANS, but the solution is linearised over the different frequencies, and thus, the LFD methods only accurately capture small displacements, just as the DLM does. It is clear that this approach, while perfectly reasonable for conventional aircraft, is not suitable for VFA.

Progress has been made towards coupling a full URANS aerodynamic solver with a nonlinear structural solver [175]. Three main philosophies can be followed when doing so: a monolithic approach, where aerodynamic and structural degrees of freedom are solved simultaneously, a strongly-coupled FSI analysis, where an iterative process between aerodynamic and structural solvers is conducted until convergence is reached for every time step, and lastly, a loosely-coupled analysis, where the aerodynamic grid is not updated at every iteration. The monolithic approach still has not been applied for complex FSI problems due to its lack of generality and overall complexity. The strongly-coupled approach has been implemented, amongst others, in ENS3DAE [155], or in the work of Silva et al. [162], where Volterra theory of nonlinear systems [103] was used to derive approximations to the behaviour of the aeroelastic model. The loosely-coupled analysis is cheaper in its evaluation [174, 176]. However, the loss of accuracy due to the lower update rate of the aerodynamics can result in longer runtimes for the same accuracy than the strongly-coupled solvers.

It is important to note that all these methods require an interpolation step between aerodynamic and structural meshes, usually performed with a 2-D or 3-D spline. Loads are transferred from the aerodynamic grid to the structural nodes, and displacements

and rotations need to be mapped onto the aerodynamic grid, and then a volume mesh deformation step needs to be conducted. Loss of mesh element quality is specially important for URANS simulations in nonlinear deformation regimes, as the volumetric mesh deformation can introduce low quality elements in the boundary layer region. In long simulations, where several deformation cycles are simulated, this loss of quality can be accumulated, resulting in even larger loss of accuracy.

Lastly, some authors have done some research quantifying the differences between predictions obtained by DLM and CFD-based aeroelasticity. [148] studied the difference on predicted gust loads for commercial aircraft, while Beaubien compared time and frequency-domain methods for flutter prediction of the AGARD 445.6 wing described in [22]. DLM is considered accurate for subsonic simulations, while it is known to lose accuracy when approaching the transonic regime [188]. Reimer's work [148] shows how the CFD-based gust loads estimations are lower than the DLM ones for both subsonic and transonic regime. However, these results are not conclusive and further research is required due to the large number of factors involved in gust-induced loads.

### 1.2.4 Computational aeroelasticity of very flexible aircraft

The aeroelastic characteristics of HALE aircraft in general and VFA in particular have been an active field of research, tackled by a growing number of researchers in the last couple of decades. While not always directly applicable to nonlinear aeroelasticity, Afonso [2] provides a useful and exhaustive survey of high aspect-ratio wings aeroelasticity with a focus on conventional commercial aviation.

One of the first papers in the field was the one published by van Schoor and von Flotow [154]. With quasi-steady modes and a 2-D strip-theory model, they obtained the quasi-steady stability properties of a very light, high aspect-ratio human-powered aircraft. Later, other developments based on similar low-order aerodynamics were published. Examples of these are [135, 136, 138], where the authors coupled composite beams with a finite-state aerodynamic model. Roughly at the same time, Drela [47] released a design and analysis tool which included a nonlinear beam model with lifting-line aerodynamics. This code was called ASWING.

Patil and Hodges developed their own code [134]. It incorporated a nonlinear beam model with a vortex-lattice and non-planar doublets for the aerodynamics. Simultaneously, Tang and Dowel [183] developed a method similarly based on nonlinear beams, but this time using a 2-D strip theory for the aerodynamics. Later on, Cesnik and his group

developed UM/NAST [24, 25] (Nonlinear Aeroelastic Simulation Toolbox), which was originally based on a strain-based beam model and finite-state aerodynamics. Recently, a UVLM code has been partially incorporated, together with added capabilities, such as propeller wake-wing interaction effects [185].

Wang was the first to efficiently couple a UVLM solver with a nonlinear structural model [203]. This represented an important step forward, as it was possible to simulate free wake and some viscous effects, which can be relevant in simulations of non conventional configurations —such as multiple-tailed flexible aircraft—, where the convected wake can interact with surfaces located downstream [126].

Lastly, SHARP (Simulation of High Aspect-Ratio Planes) [133], and more recently SHARPy (Simulation of High Aspect-Ratio aeroplanes in Python) [39], is the code developed at Imperial College London. It uses a displacement and rotations-based geometrically-exact composite beam model [125, 168, 70] with multibody constraints, together with a UVLM for the aerodynamics [124, 169]. The FSI coupling is formulated as a strongly-coupled problem with the possibility of conducting loosely-coupled simulations as well. SHARPy has nonlinear flight dynamics simulation support, together with linearisation and model reduction capabilities. At the time of writing this, SHARPy is the only nonlinear aeroelasticity toolbox released freely under an Open Source license. The work presented in this thesis has been carried out entirely with SHARPy, unless explicitly indicated (usually for validation purposes).

All the authors cited previously in this section were focusing on low and medium fidelity aerodynamic models. In a parallel effort, high-fidelity CFD was being investigated as a potential aerodynamic solver for coupled simulations. García and Guruswamy [59] coupled a Navier-Stokes solver to a nonlinear beam solver, and Smith, Patil and Hodges [175] used the higher fidelity of the aerodynamics solver, coupled with a geometrically nonlinear beam model in order to assess the effect of nonlinearities in the aeroelastic analysis. Palacios and Cesnik [132] coupled a quasi-three-dimensional structural solver with a Navies-Stokes aerodynamic model. This work highlighted the importance of sectional deformations in the location of the shock wake at transonic regimes. Lastly, Hallisy [67] coupled an unstructured 3-D CFD solver with the quasi-three-dimensional model of Palacios.

This last section has provided an overview of the possible choices for computational aeroelasticity of very flexible aircraft. Given the scope of this work, that deals with flight speeds in the low subsonic regime and generally attached flow, potential methods are better suited due to their much lower cost and ease of coupling. In addition, the

UVLM formulation used in this work predicts the induced component of the drag, which is arguably one of the main design targets for HAPS configurations. Viscous drag can be easily added as a correction.

## 1.3 Atmospheric boundary layer modelling for linear and nonlinear aeroelasticity

Atmospheric turbulence and its effects on aircraft flight have been a major concern of the aerospace sciences for more than a century [79]. The efforts to characterise the impact of atmospheric turbulence on aircraft loads can be traced back to the early development and consolidation of discrete gust and continuous turbulence models [86, 116, 213, 49, 76, 72]. Today, airworthiness certification procedures (amongst them [53, 52] for manned air vehicles and [127] for unmanned platforms) have adapted guidelines and methodologies derived from both that early research and the accumulated experience of decades of in-service flight data. Hoblit [72] provides a good overview of how these requirements were created and shaped in the 1950s for discrete gusts and 1970s for continuous turbulence models. However, the very recent development of high-altitude pseudo satellites (HAPS) for an increasingly wide range of applications in remote sensing and communications [189, 212] is challenging the conventional understanding of the interactions between (nonlinear) aeroelasticity and atmospheric turbulence. The extreme operational demands and requirements of HAPS have resulted in more slender and much lighter platforms which inevitably renders them more flexible. This is true for example for stratospheric solar high-aspect ratio aircraft (NASA Helios, Boeing Oddyseus, Facebook Aquila, Airbus Zephyr, etc. — see Section 1.2.1) [65]. Poor understanding of the low-altitude response of such very flexible aircraft (VFA), together with their low controllability and slow characteristic times has previously led to poorly controlled flight dynamics and even vehicle loss [130]. Flight in low altitude turbulence response is a challenging case due to the anisotropy of the flow, and the typical size of the eddies, which together with the low flight speed of these aircraft, mean that one-dimensional turbulence assumptions need to be revisited. Gaining physical insight into the aeroelastic interactions while quantifying the uncertainty of flight dynamics is considered key in further developing these technologies.

### 1.3.1 Atmospheric boundary layer field generation

In simulating atmospheric turbulence, two approaches can be generally considered. The first approach uses stochastic formulations to generate wind fields based on first and second-order statistics (spectral/co-spectral). Stochastic generation of wind fields has been an active area of research over the last decades [194, 105, 106]. Aicraft continuous loads due to atmospheric disturbances have mostly been computed using von Kármán spectrum [44], while in the early years of frequency-domain for aeroelasticity analysis, the Dryden spectrum [100] was used due to its easier implementation, as it is given in a rational transfer function shape and requires no simplification. It is also used to test the effect of turbulence in flight dynamics applications [115, 153]. Other 1-D models of atmospheric turbulence have been proposed, in particular for flight dynamics applications. While they will not be discussed here, the most common ones are summarised in [199]. For applications in wind engineering [82, 60], the most common stochastic model is that of Kaimal [88, 194]. The second technique on the other hand, utilises detailed numerical simulations to resolve the fluid flow equations (or a parametrisation of them) for a given temporal and spatial scale [140]. Such a popular method is large-eddy simulation (LES). The idea underpinning LES lies in resolving the large energetic structures of the flow while modelling the smaller scales (SGS, subgrid scales). LES produces three-dimensional wind fields for a given computational domain and its quality depends on the resolution and modelling parameters. To this end, LES not only does it resolve all relevant turbulent scales, it provides (when properly designed) excellent agreement with theoretical model such as the Monin-Obhukov similarity theory (MOST) [121] and experimental data [68].

#### 1.3.1.1 Spectral turbulence generation methods for aeronautics

Around the 1960s, the first attempts to develop a correct general set of power-spectral gust models were undertaken. Before then, only discrete sharp-edge or 1–cos gusts were used to approximate the structural loads due to atmospheric disturbances. Nowadays, CS-25 [52] (and FAR-25, by extension) Acceptable Means of Compliance explicitly encourage the use of the von Kármán spectrum as the accepted way of certifying commercial aircraft. It is interesting to point out that small and commuter aircraft, which have to comply with another certification specification (CS-23 [53]), are not subject to continuous turbulence requirements, and only a small set of 1–cos discrete gusts are required.

Stochastic turbulence models, such as the von Kármán one, describe a stochastic process with a number of assumptions [51]:

*Gaussian*: the probability distribution of disturbance velocities is Gaussian

*Stationary*: process statistics are constant in time

*Homogeneous*: the process statistics do not depend on the location they are sampled at

*Ergodic*: the process statistics can be deduced form a single, long enough, sample of the process

*Frozen*: the velocity field does not vary in time, but it does so in space

*Isotropic*: the process statistics do not depend on the orientation of the axes. It is important to note that some spectral methods include some variation of the length scales with altitude.

Ever since the early stages of the spectral methods use for loads prediction, it has been known that real-life turbulence is not exactly Gaussian, nor it complies with the above list of assumptions. A good discussion on this can be found in [72]. From this reference, the most relevant discrepancies for the synthetic flow fields in comparison with real-life turbulence of interest for this work are:

the lack of intermittency and strong vortex patterns,

non-*ergodicity* character of real turbulence,

lower peak-to-rms response,

and lastly, even for almost-Gaussian real-life turbulence, the probability distribution of the velocity differences depart substantially from the assumed distribution.

Despite this, the applicability of spectral turbulence generation methods has been proven beyond doubt for commercial aircraft application. The higher flight velocity of these aircraft when compared with HALE configurations, together with their comparable dimensions, requires an additional analysis. For example, the non-ergodicity of the synthetic fields can play a role in the response, given how much slower HALE aircraft are.

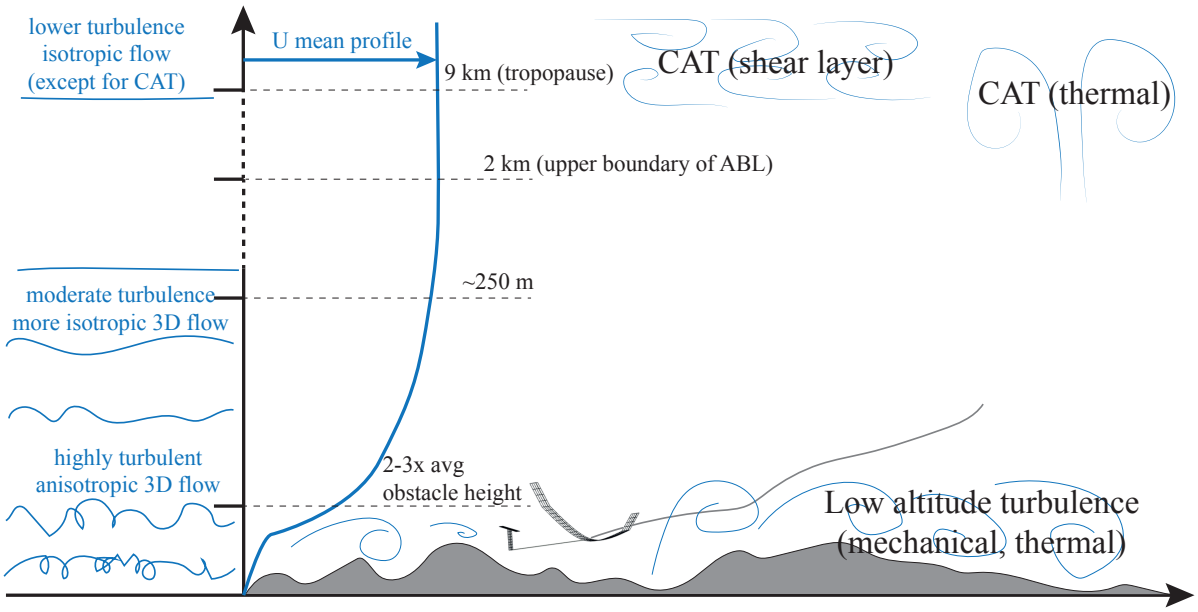



**Fig. 1.5** Schematic representation of the atmospheric turbulence experienced by the flexible aircraft.

### 1.3.2 Atmospheric boundary layer aeroelastic effects

Continuous turbulence design criteria are often based on the von Kármán spectrum with a few empirical observations. Employing such approach to study atmospheric turbulence interacting with very flexible aircraft aeroelasticity poses two major issues. First, it has been already said that the von Kármán spectrum is derived from isotropic turbulence arguments, a condition which is rarely met in atmospheric turbulence, and never at low altitudes. Some military standards (MIL-A-8861) have tried to include this effect in the classic von Kármán spectrum by adding a turbulence characteristic length ($L$) dependency with the altitude. In the lower part of the planetary boundary layer (surface layer), turbulence is strongly modulated by the presence of an aerodynamically rough surface and heat flux due to solar radiation. These effects combined together can significantly alter the flow dynamics and therefore deviate from the isotropic turbulence case implied by the von Kármán-like turbulence that is usually found in higher altitudes (e.g. Clear-Air Turbulence, CAT). The low-altitude operating environment of a generic VFA is shown schematically in figure 1.5. The near-surface turbulence coherent structures (e.g. high/low-speed streaks) will dominate the flight dynamics of such vehicle.

The second shortcoming relates to the coupling between these large-scale turbulence structures (e.g. eddies) with the dynamics of very flexible aircraft. As high wing deformations are expected under normal trim conditions, aircraft are prone to dynamic

instabilities of significant magnitude [2, 130]. In particular, high-aspect ratio wings are expected to present a coupling between rigid-body and elastic motions since the frequency of the elastic modes are closer to those of the flight dynamics modes. This strong coupling has been well-documented in, e.g., [153] and [191] for general aircraft aeroelasticity and, more specifically for high-aspect ratio aircraft, in [137, 182, 145, 178], among others. To this day, a limited number of published experimental studies on high-aspect ratio, flexible wings exists ([209, 31]). They have shown however, that very flexible wings have a significantly different gust response characteristics under different deformations and load conditions ([102, 30]). Inherently, additional analysis is required in order to gain insight into the complex interactions between wind gusts and aircraft aeroelasticity. This particular study focuses on investigating the effect of low-altitude turbulence on very flexible high-aspect ratio aircraft by first conducting a sensitivity analysis of atmospheric turbulence generation methods including the industry-standard von Kármán model as well as demonstrating the differences in the dynamic aeroelastic response between very flexible and stiff aircraft.

The computational power and the versatility in methods and implementation required to simulate the response of an aircraft immersed in a realistic 3-D turbulent field have delayed the analysis of this kind of problems until recently, and these new simulations are based on rigid aircraft, 6-DOF approximations to flight dynamics that would not be accurate for VFA-related cases. Work by Duraisamy [36] studies the flight dynamics under a controller of a quadcopter immersed in a LES boundary layer with snapshots captured every 12.1 s, amongst other wind models. The rigid quadcopter model is based on a blade-element momentum theory solver with a nonlinear 6 degrees of freedom controller for the flight dynamics. The work presented here introduces higher-fidelity turbulence fields, with better spatial and temporal resolution, a more advanced flight dynamics solver that integrates geometrically nonlinear dynamics, and a more sophisticated coupled aerodynamics solver. This is, to the knowledge of the author, a new development in the field.

## 1.4 Research questions of this thesis

HALE configurations face a different set of challenges than the ones identified in conventional aircraft aeroelasticity. Identifying, modelling and solving these is an ongoing effort in which analysis tools need to be constantly revisited as the engineering ambition in those platforms increases. For example, up until now, the usual approach to HALE aircraft gust

loads analysis was to use simplified, partly nonlinear formulations to size and certify the aircraft using the industry-standard cases, such as the von Kármán continuous turbulence and a 1-cos span-constant gust. We aim to revisit the design/analysis challenges that occur at low altitudes.

Additionally, no other nonlinear aeroelasticity toolbox available is capable of simulating complex highly dynamical and nonlinear manoeuvres. Procedures such as a catapult-assisted take-off could only be simulated using quasi-steady approximations and linear solvers.

This work aims to fill the gap in knowledge and capabilities that exists nowadays in the context of highly flexible HALE aircraft:

1. A study of the suitability of conventional sizing rules and certification procedures for highly flexible, span-loaded, low-speed aircraft is necessary. This work is part of an effort to explore and develop nonlinear aeroelastic tools suitable for these configurations and suitable for industrial applications. In this same context, how do the worst case scenarios look like, and which tools are necessary for their simulation?

2. What is an appropriate modelling strategy to estimate structural loads under low-altitude winds on slowly-moving, high aspect-ratio aircraft?

3. When simulating complex manoeuvres of very flexible aircraft, which numerical parameters affect the solution, which ones accelerate the convergence of the method and which are less relevant than in conventional aircraft aeroelasticity? What specific considerations need to be taken into account in order to ensure the correctness of the results?

4. Once suitable aeroelastic simulation methods have been identified and developed, how can these be used in order to optimise complex nonlinear manoeuvres in a versatile and inexpensive way? In addition, which optimisation strategy is most suitable for this kind of problem, where function calls are expensive and the dimensionality of the problem is relatively low?

## 1.5 Dissertation outline

- Chapter 1 has provided an overview of the aim of this work, together with a literature review covering the most important topics for this dissertation, including state of the HALE industry, nonlinear aeroelasticity and turbulence field generation methods.

The chapter finishes enumerating the research questions of this dissertation, together with the list of original contributions that were published as a result from this effort.

- Chapter 2 introduces the numerical methods used in this work, from the geometrically nonlinear structural solver up until the aeroelastic coupling scheme. Extended discussions concerning modelling strategies for HALE-like platforms are included, together with additional learnt lessons for more complex, although common problems, such as longitudinal trim routines for VFA.

- Chapter 3 compiles a number of validation or interesting cases and the aircraft models description that are used throughout the dissertation. The introduction of the validation cases is used as starting point for several discussions on modelling choices and parameter selection that were not covered in past literature or the previous chapter. These cases cover geometrically nonlinear structural solutions, UVLM validation, and static nonlinear solutions for two different aircraft.

- Chapter 4 deals with the atmospheric turbulence modelling impact on VFA aircraft aeroelasticity. Turbulence field generation and particular challenges arising from non linear aeroelastic simulations immersed in such fields are discussed. Potential bottlenecks are identified and solutions concerning model and implementation approaches are identified. The turbulence model used to generate the flow field is proven to have a noticeable impact in the structural loads estimation for high aspect-ratio, flexible aircraft. The impact of the flexibility effects is also studied for the same aircraft model.

- Chapter 5 defines and parametrises a complex nonlinear manoeuvre: a Catapult-Assisted Take Off of a VFA. A baseline launch, a set of constraints, and a cost function related to its suitability from a ground clearance and wing root loads point of view are defined. A custom-developed hybrid global/local optimiser based on Bayesian Optimisation and gradient-less local methods is described. The Catapult-Assisted Take Off is used as an application example of this optimiser. In addition, specific numerical considerations are identified and best practices are stated for the CATO simulation, as well as the optimisation process in the context of expensive, nonlinear cost functions.

- Chapter 6 summarises the results obtained during this work. Conclusions concerning every part of the discussion are presented.

## 1.6 Original contributions of this work

- Fully nonlinear time-domain simulations of very flexible aircraft immersed in low altitude, time-varying 3-D atmospheric boundary layer have been conducted. Bottlenecks and potential solutions are identified and described. The continuous turbulence response for HALE-like aircraft has been analysed for three turbulence generation methods in increasing order of complexity and dimensionality of its coherence: von Kármán (1-D coherence), Kaimal (2-D coherence) and LES (3-D coherence). Load envelopes, rigid-body motion and structural deflections are presented and compared.

- The effect of the spatial discretisation of the aerodynamic grid is discussed for longitudinal and lateral dynamics. The model used for this is a representative, real life HALE aircraft: the X-HALE, a real-life, flexible model built at University of Michigan (UM) for nonlinear aeroelasticity research. Results are compared with UM/NAST, the in-house code at UM. The differences between strongly coupled and weakly coupled and strain vs. displacements-based structural formulations are analysed.

- With the challenges of using a nonlinear aeroelastic simulation as cost function in mind, a hybrid (global/local) custom optimisation process based on Bayesian Optimisation has been designed.

- A framework for the simulation of a catapult-assisted take off has been designed and applied to the design of the manoeuvre. Nonlinear and unsteady contributions to the dynamics of the problem have been identified. The elasticity effects have been quantified, as well as the robustness of the solution when the conditions vary. This has been introduced in the custom optimisation routine in order to determine the optimal procedure.

- All the above have been made available in a unique Open Source software: SHARPy. The core of this code supports the analysis of steady and unsteady nonlinear aeroelastic response of aircraft and wind turbines, in open and closed loop, to any kind of atmospheric disturbance. The modular architecture of it has enabled the integration of work from other members of the Loads Control and Aeroelastics laboratory. The most relevant being: a linear and asymptotic stability aeroelastic analysis module, extended support for wind turbines and multibody constraints,

and model order reduction applied to aeroelastic systems. All these features are available, automatically tested and fully documented. SHARPy is now being used in industry and academia. Additionally, several publications [7, 122] have used this software as their software basis, simultaneously contributing to the code.

## 1.7 Publications based on this work

- Del Carre, A. and Palacios, R., 2019. Efficient Time-Domain Simulations in Nonlinear Aeroelasticity. *AIAA Scitech 2019 Forum, San Diego, CA*.
- Deskos, G., del Carre, A. and Palacios, R., 2020. Assessment of low-altitude atmospheric turbulence models for aircraft aeroelasticity. *Journal of Fluids and Structures 2020. Vol. 95*
- Del Carre, A., Teixeira, P., Cesnik, C. E. S., and Palacios, R., Nonlinear response of a very flexible aircraft under lateral gust, 2019. *IFASD Forum, Savannah, Georgia, 2019.*
- Del Carre, A., Muñoz-Simón, A., Goizueta, N., and Palacios, R., 2019. SHARPy: A dynamic aeroelastic simulation toolbox for very flexible aircraft and wind turbines. *Journal of Open Source Software, 4(44):1885.*
- Del Carre, A. and Palacios, R., Realistic turbulence effects in low altitude dynamics of very flexible aircraft, 2020. *AIAA Scitech 2020 Forum, Orlando, FL*
- Del Carre, A. and Palacios, R., Simulation and optimization of takeoff maneuvers of very flexible aircraft, 2020. *AIAA Journal of Aircraft 2020*, Accepted.

# Chapter 2

# Numerical methods

This chapter outlines the main features of numerical methods used in the rest of the dissertation, including a critical assessment of the capabilities and domain of application of each method.

Section 1.1 has justified the choice of structural and aerodynamic solver. In short, given the expected 1) large structural deformations, 2) strong rigid-body motion-elastic degrees of freedom coupling, 3) three-dimensional effects and induced drag prediction requirements, and 4) potential importance of wake convection, the range of choice is relatively narrow. Concerning the structural model, geometrically-linear and quasi-linear solvers are discarded, however, structural linearities in the constitutive equations is a reasonable assumption. Due to the slender character of the airframe, a 1-D approximation to the structural is acceptable, as long as stiffness cross-coupling terms are supported. The aerodynamic solver needs to support the arbitrarily large deformations of the lifting surfaces, be able to simulate the free wake dynamics and natively consider three-dimensional effects with no addition user input.

The numerical methods used in this work are a Geometrically-Exact Composite Beam (GECB) structural model and an Unsteady Vortex-Lattice Method (UVLM) aerodynamic model.

## 2.1 Geometrically-Exact Composite Beam

The structural solver presented here is a finite element beam model formulated in nodal displacements and rotations with quadratic elements [62, 168]. As explained previously, this formulation is particularly suitable for nonlinear aeroelasticity due to the

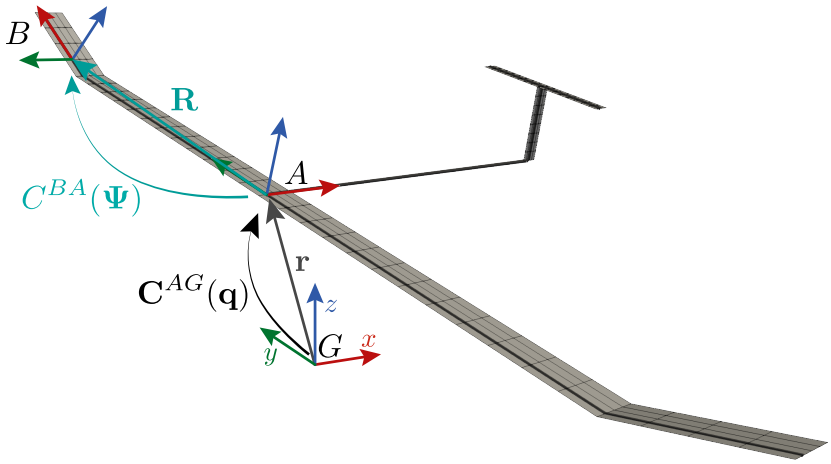



**Fig. 2.1** Different frames of reference used in the description of the problem.

native nonlinearity and follower force support, as well as the ease of coupling with the aerodynamic model and the boundary condition implementation. In addition, the solver supports the inclusion of multibody constraints through Lagrange Multipliers (section 2.1.3).

Large deformations can result in large rotations, which are parametrised using the Cartesian Rotation Vector (CRV) description. Among its advantages we can underline its memory efficiency, with three degrees of freedom per node, and the easy support of linearisation techniques around an initially-curved beam geometry. It has to be noted that this formulation on displacements and CRV rotations does not preserve the objectivity of the rotations. However, work conducted by Simpson [168] showed that using quadratic elements and considering element sizes that are reasonable for the aerodynamic model as well, the error due to this is negligible.

### 2.1.1 Beam kinematics

Figure 2.1 shows the frames of reference involved in the description of the structural problem, and their relationship and parametrisations. The description of the deformation of the structure is given in a body-fixed, moving frame of reference $A$. Its kinematics with respect to the inertial frame, $G$, are defined by the instantaneous translational velocity $\boldsymbol{v}_A(t)$ and angular velocity $\boldsymbol{\omega}_A(t)$ of its origin. Subindices are used to express that vector magnitudes are projected onto their components on a particular reference frame. We define the aggregated vector of rigid-body velocities $\boldsymbol{\beta} = \left\{\boldsymbol{v}_A^\top \boldsymbol{\omega}_A^\top\right\}^\top$, where an uppercase subindex indicates the frame of reference in which the variable is evaluated.

The orientation of $A$ with respect to $G$ is parametrised using quaternions $\boldsymbol{q}(t) = \left[q_0(t)\ \boldsymbol{q}_v(t)^T\right]^T$ and alternatively described as $\boldsymbol{C}^{GA}(\boldsymbol{q})$. The use of quaternions for this rotation allows an efficient and accurate description of the possibly large rigid-body motion of the flexible body while avoiding accumulated errors due to consecutive finite rotations and gimbal lock effects. In addition, the rotation of $A$ in inertial FoR can be easily tracked using the *attitude propagation* equations:

$$\begin{bmatrix} \dot{q}_0 \\ \dot{\boldsymbol{q}}_v \end{bmatrix} = -\frac{1}{2} \begin{bmatrix} 0 & \boldsymbol{\omega}_A^T \\ \boldsymbol{\omega}_A & \tilde{\boldsymbol{\omega}}_A \end{bmatrix} \begin{bmatrix} q_0 \\ \boldsymbol{q}_v, \end{bmatrix} \tag{2.1}$$

The inertial coordinates of a point located in a beam cross-section can be obtained following Fig. 2.1. With the coordinates of the point with respect to the beam elastic axis in material $(B)$ frame noted as $\boldsymbol{\xi}_B$, the inertial coordinates, $\boldsymbol{X}_G$ are given as

$$\boldsymbol{X}_G = \boldsymbol{r}_G + \boldsymbol{C}^{GA}\boldsymbol{R}_A + \boldsymbol{C}^{GA}\boldsymbol{C}^{AB}\boldsymbol{\xi}_B, \tag{2.2}$$

with the variables in here defined in figure 2.1.

The local orientation of the beam cross-sections is given as the associated material frame of reference $B$ orientation. The rotation from body-attached frame, $A$, to material, $B$, is parametrised using the Cartesian Rotation Vector description: $\boldsymbol{\Psi}(s,t)$. The CRV is based on Euler's rotation theorem, that states that any 3-D rotation of a rigid body around a fixed point can be reduced to a single rotation of magnitude $\phi$ around a fixed axis, $\hat{\boldsymbol{n}}$. From this, the CRV can be defined as $\boldsymbol{\Psi} = \phi\hat{\boldsymbol{n}}$, and then $\boldsymbol{C}^{BA}(s,t) = \boldsymbol{C}\left(\boldsymbol{\Psi}(s,t)\right)$ can be computed as

$$\boldsymbol{C}(\boldsymbol{\Psi}) = \boldsymbol{I} + \frac{\sin\phi}{\phi}\tilde{\boldsymbol{\Psi}} + \frac{1-\cos\phi}{\phi^2}\tilde{\boldsymbol{\Psi}}^2 = \sum_{k=0}^{\infty}\frac{1}{k!}\tilde{\boldsymbol{\Psi}}^k, \tag{2.3}$$

with $\tilde{\bullet}$ indicating the skew-symmetric matrix operator. The first expression for the rotation matrix presents a division by zero in $\phi \to 0$, which results in unstable behaviour, or errors when implemented in a computer. In contrast, the last term in equation 2.3 is the exponential map representation of $\boldsymbol{C}(\boldsymbol{\Psi})$ and can be used for small angles. It is easy to see that this representation eliminates the singularity for $\phi = 0$. The problem of the exponential map representation is that the computation of the exponential of a matrix is very costly and it is still an open topic of research (the most common way is to use Padé polynomials to approximate the solution [120]), and because this function is called

repeatedly during the solution of the problem, its cost accumulates very rapidly. For this work, the first expression in (2.3) is used for large $|\phi| > \epsilon_\phi$, where $\epsilon_\phi$ is a given threshold. In the $|\phi| \leq \epsilon_\phi$ case, a linearised version of the expression, based on the truncated series, is used:

$$\boldsymbol{C}(\boldsymbol{\Psi}) = \boldsymbol{I} + \tilde{\boldsymbol{\Psi}} + \frac{1}{2}\tilde{\boldsymbol{\Psi}}^2. \tag{2.4}$$

A too high value for $\epsilon_\phi$ causes loss of accuracy due to the low order component of the error in equations (2.4) (and (2.7) later on). On the other hand, if $\epsilon_\phi$ is too small, some numerical noise (or floating-point exceptions) can appear due to floating-point arithmetic.

In order to obtain expressions for the local strain and velocity in terms of the derivatives of the CRV, a perturbation of the rotation matrix, $\boldsymbol{C}(\boldsymbol{\Psi})$ is evaluated. The matrix of infinitesimal rotation of the local material frame $B$ is given by

$$\widetilde{\delta\boldsymbol{\Phi}_B} = \boldsymbol{C}\delta\boldsymbol{C}^T = \boldsymbol{T}\widetilde{\delta\boldsymbol{\Psi}}, \tag{2.5}$$

where the rotational operator $\boldsymbol{T}$ is derived in Ref. [62], and it can be written as

$$\boldsymbol{T}(\boldsymbol{\Psi}) = \boldsymbol{I} + \frac{\cos\phi - 1}{\phi^2}\tilde{\boldsymbol{\Psi}} + \frac{1 - \frac{\sin\phi}{\phi}}{\phi^2}\tilde{\boldsymbol{\Psi}}\tilde{\boldsymbol{\Psi}}. \tag{2.6}$$

In the same manner as with the rotation matrix in equation (2.3), when $|\psi| \leq \epsilon_\phi$, a truncated series expression is used:

$$\boldsymbol{T}(\boldsymbol{\Psi}) = \boldsymbol{I} - \frac{1}{2}\tilde{\boldsymbol{\Psi}} + \frac{1}{6}\tilde{\boldsymbol{\Psi}}^2. \tag{2.7}$$

It is easy to see that with $\phi \to 0$, both the rotational and the tangential operators tend to $\boldsymbol{I}$.

### 2.1.2 Beam dynamics

This section is based on the work by Palacios, Murua and Cook [133]. The geometrically exact composite beam model is built from the equations of motion of a curved, nonlinear beam of arc-length $\ell$ for the time interval $[t_1, t_2]$. Using the subindex of every term as an indication of the frame of reference in which the quantity is expressed, Hamilton's principle in its strong form applied on an homogenised composite structure can be written

as [69]:

$$\begin{aligned}\left(\frac{\mathrm{d}}{\mathrm{d}t}+\tilde{\boldsymbol{\Omega}}_B\right)\boldsymbol{P}_B &= \left(\frac{\mathrm{d}}{\mathrm{d}s}\tilde{\boldsymbol{K}}_B\right)\boldsymbol{F}_B+\boldsymbol{f}_B\\ \tilde{\boldsymbol{V}}\boldsymbol{P}_B+\left(\frac{\mathrm{d}}{\mathrm{d}t}+\tilde{\boldsymbol{\Omega}}_B\right)\boldsymbol{H}_B &= \left(\frac{\mathrm{d}}{\mathrm{d}s}\tilde{\boldsymbol{K}}_B\right)\boldsymbol{M}_B+(\tilde{\boldsymbol{e}}_1+\tilde{\boldsymbol{\gamma}})\cdot\boldsymbol{F}_B+\boldsymbol{m}_B.\end{aligned} \tag{2.8}$$

where $\boldsymbol{V}$ is the beam local translational velocity, $\boldsymbol{P}$ is the beam local translational momenta, $\boldsymbol{\Omega}$ and $\boldsymbol{H}$ are the beam local rotational velocities and momenta, respectively. Next, $\boldsymbol{\gamma}$ is the beam local force strain and $\boldsymbol{F}$ is the beam sectional internal forces vector. Correspondingly, $\boldsymbol{\kappa}$ is the beam local moment strain and $\boldsymbol{M}$ is the local sectional internal moments. $\boldsymbol{e}_1$ is the $[1\ 0\ 0]^T$ vector. Lastly, $\boldsymbol{f}$ is the local external forces per unit length, and $\boldsymbol{m}$ is the local external moments per unit length. All vectors are expressed in their components in the local material frame, $B$.

The constitutive equations are assumed linear:

$$\begin{aligned}\boldsymbol{P}_B &= \bar{m}\boldsymbol{V}_B-\bar{m}\tilde{\boldsymbol{\xi}}_{cg}\boldsymbol{\Omega}_B, \quad \boldsymbol{\gamma}=c_{\gamma f}\boldsymbol{F}_B+c_{\gamma m}\boldsymbol{M}_B,\\ \boldsymbol{H}_B &= \bar{m}\tilde{\boldsymbol{\xi}}_{cg}\boldsymbol{V}_B+\boldsymbol{J}\boldsymbol{\Omega}_B, \quad \boldsymbol{\kappa}=c_{\kappa f}\boldsymbol{F}_B+c_{\kappa m}\boldsymbol{M}_B,\end{aligned} \tag{2.9}$$

where $c$ are the linear cross-sectional stiffness coefficients, $\bar{m}$ is the mass per unit length, $\boldsymbol{\gamma}$ and $\boldsymbol{\kappa}$ are the strains and curvatures, respectively, and $\boldsymbol{\xi}_{cg}$ are the cross-sectional coordinates of the sectional centre of mass with respect to the elastic axis.

Equations (2.8) and (2.9) are solved together with the nonlinear velocity-displacement kinematic relations:

$$\begin{aligned}\boldsymbol{V}_B &= \boldsymbol{C}^{BA}\left(\dot{\boldsymbol{R}}_A+\tilde{\boldsymbol{\omega}}_A\boldsymbol{R}_A+\boldsymbol{v}_A\right),\\ \tilde{\boldsymbol{\Omega}}_B &= \boldsymbol{C}^{BA}\left(\dot{\boldsymbol{C}}^{AB}+\tilde{\boldsymbol{\omega}}_A\boldsymbol{C}^{AB}\right).\end{aligned} \tag{2.10}$$

Next, from the deformation and rotation data, the strain and curvature vectors can be computed as:

$$\begin{aligned}\boldsymbol{\gamma}(s,t) &= \boldsymbol{C}\left(\boldsymbol{\Psi}(s,t)\right)\boldsymbol{R}'_A(s,t)-\boldsymbol{C}\left(\boldsymbol{\Psi}(s,0)\right)\boldsymbol{R}'_A(s,0)\\ \boldsymbol{\kappa}(s,t) &= \boldsymbol{T}\left(\boldsymbol{\Psi}(s,t)\right)\boldsymbol{\Psi}'(s,t)-\boldsymbol{T}\left(\boldsymbol{\Psi}(s,0)\right)\boldsymbol{\Psi}'(s,0)\end{aligned} \tag{2.11}$$

Choosing the CRV as the rotation parametrisation and combining equations (2.8), (2.9), (2.10) and (2.11), a closed system of equations is obtained. This system is then discretised with quadratic 1-D finite elements. Being $\boldsymbol{\eta}$ the state variable containing the displacements and rotations and $\boldsymbol{\mathcal{N}}$ the shape functions, we have: $\boldsymbol{\eta}=\boldsymbol{\mathcal{N}}\bar{\boldsymbol{\eta}}$. The

rewritten Hamilton's principle results:

$$\boldsymbol{\mathcal{M}}_\eta(\bar{\boldsymbol{\eta}})\ddot{\bar{\boldsymbol{\eta}}} + \boldsymbol{m}_a(\bar{\boldsymbol{\eta}}) \begin{bmatrix} \dot{\boldsymbol{v}}_A \\ \dot{\boldsymbol{\omega}}_A \end{bmatrix} + \boldsymbol{\mathcal{F}}_{\text{gyr}}(\bar{\boldsymbol{\eta}}, \dot{\bar{\boldsymbol{\eta}}}, \boldsymbol{v}_A, \boldsymbol{\omega}_A) + \boldsymbol{\mathcal{F}}_{\text{stif}}(\bar{\boldsymbol{\eta}}) = \boldsymbol{\mathcal{F}}_{\text{ext}}, \tag{2.12}$$

where $\boldsymbol{\mathcal{M}}_\eta$ is the discrete mass matrix –which may include non-structural mass–, and the discrete gyroscopic, stiffness and external generalised forces are represented by $\boldsymbol{\mathcal{F}}_{\text{gyr}}$, $\boldsymbol{\mathcal{F}}_{\text{stif}}$, and $\boldsymbol{\mathcal{F}}_{\text{ext}}$, respectively.

It is important to note that the applied forces are given in material ($B$) frame of reference, and as such, they will rotate and move with the beam node they are associated to. The orientation and position is updated at every structural iteration. This is one of the reasons the GECB as it is implemented here performs so well in a nonlinear aeroelasticity context.

The rigid-body degrees of freedom are included by augmenting the system with the $\boldsymbol{\beta}$ vector, which includes the $A$ frame of reference linear and angular velocities and orientation. The resulting system is

$$\boldsymbol{\mathcal{M}}(\boldsymbol{\eta}) \begin{bmatrix} \ddot{\boldsymbol{\eta}} \\ \dot{\boldsymbol{\beta}} \end{bmatrix} + \begin{bmatrix} \boldsymbol{\mathcal{Q}}^S_{\text{gyr}} \\ \boldsymbol{\mathcal{Q}}^R_{\text{gyr}} \end{bmatrix} + \begin{bmatrix} \boldsymbol{\mathcal{Q}}^S_{\text{stif}} \\ 0 \end{bmatrix} = \begin{bmatrix} \boldsymbol{\mathcal{Q}}^S_{\text{ext}} \\ \boldsymbol{\mathcal{Q}}^R_{\text{ext}} \end{bmatrix} \tag{2.13}$$

This equation is then solved iteratively using a fixed-point iteration scheme. Time integration is carried out with an explicit Newmark-$\beta$ [1] scheme formulated in incremental form [160]. Being $\boldsymbol{\eta}^k$ the state vector at the $k$-th iteration and $\Delta\boldsymbol{\eta}^k$ the increment from the Newmark-$\beta$ iterative process, the next iteration state vector is given by:

$$\boldsymbol{\eta}^{k+1} = \boldsymbol{\eta}^k + \Delta\boldsymbol{\eta}^{k+1}. \tag{2.14}$$

The convergence residual is defined as $\epsilon^k = \|\Delta\boldsymbol{\eta}^k\|/\|\Delta\boldsymbol{\eta}^0\|$. Additionally, an absolute convergence threshold is implemented for cases where the initial state is close to the converged solution.

The Newmark-$\beta$ scheme introduces a numerical damping parameter that damps the high-frequency modes (which potentially are more destabilising). It has been observed that as long as the damping parameter is below a certain threshold, there is no effect in the solution for the frequencies of interest, while the destabilising numerical behaviour is appreciably damped. Past a critical value that depends on the structure, the relevant structural dynamics of the solution become artificially damped.

In cases where the numerical solution is unstable due to the large deformations, two methods have been devised to improve the convergence of the method. In both static and dynamic simulations, a relaxation parameter can be applied to the $\Delta\boldsymbol{\eta}^{k+1}$ term of equation (2.14) so that it results

$$\boldsymbol{\eta}^{k+1}\Big|_{\text{relax}} = \boldsymbol{\eta}^{k} + (1-\omega)\Delta\boldsymbol{\eta}^{k+1}, \quad \omega \in [0,1). \tag{2.15}$$

In the case of static coupled simulations where the static deflection is large enough to cause divergence in the structural solver, the best solution is to introduce the aerodynamic, gravity and applied loads in a ramp so that $N_{\text{load steps}}$ coupled cases, and the converged solution of each is used as initial condition for the next. While this slows down the static solution process, it is valuable in cases where the flight conditions might not be the nominal $g = 1$ case. For example, when running an aeroelatic trim routine[1], the initial condition can be far from the aeroelastically trimmed solution, and this can be made even more challenging if by chance, the epsilon used for the finite differences runs the static solver in even more different conditions. For these cases, the load stepping procedure can improve the robustness of the solver in exchange for longer runtimes.

### 2.1.3 Multibody constraints

Additional constraints may be needed on the flexible aircraft equations when modelling multibody problems (e.g., hinged winglets) or for problems with forced displacements (as in the catapult-assisted takeoff, presented in Chapter 5.1). This work has not been conducted by the author, the author being Arturo Muñoz Simón, from the Loads Control and Aeroelasticity Lab at Imperial College London. It is included here for completeness sake.

Different approaches have been studied to solve this kind of systems [164]. Amongst them, the *Lagrange Multipliers* (LM) method creates a new Lagrangian of the system that includes $M$ unknown multipliers (the vector of these noted $\boldsymbol{\lambda}$) that are chosen such that the principle of virtual work is enforced. This way, a system of $N + M$ equations and unknowns is defined. The implemented variant of the LM method is the *Augmented Lagrangian* approach to increase numerical robustness.

[1] In this work, the use of *aeroelastically trimmed* configuration implies the aeroelastic solution where the aerodynamic and the structural forces are in equilibrium, not the more common trimmed solution, where lift equals weight and drag equals thrust.

It is modification of the original LM where a penalty function is included on top of the constraint. This allows to accurately enforce the constraints.

From a mathematical point of view, the Lagrange Multipliers are a constrained minimisation problem. Given the definition of the action of a mechanical system

$$\mathcal{S}[\boldsymbol{q}] = \int_{t_1}^{t_2} \mathcal{L}\left(\boldsymbol{q}(t), \dot{\boldsymbol{q}}(t), t\right) \mathrm{d}t \tag{2.16}$$

where $\mathcal{L}\left(\boldsymbol{q}, \dot{\boldsymbol{q}}, t\right)$ is the Lagrangian function. The evolution of a system of action $\mathcal{S}$ subject to the constraints $\boldsymbol{\Phi}$ can be put as

$$\begin{cases} \min \mathcal{S}(\boldsymbol{q}, \dot{\boldsymbol{q}}, t) \\ \text{subject to } \boldsymbol{\Phi}(\boldsymbol{q}, \dot{\boldsymbol{q}}, t) = 0 \end{cases} . \tag{2.17}$$

The constraints $\boldsymbol{\Phi}$ can be classified as *holonomic*—they are a function of the generalized coordinates and potentially time—, and *non-holonomic*. Non-holonomic constraints can be *bilateral* when they express equality or *unilaterial* when they are an inequality.

The multibody solver in SHARPy implements the augmented Lagrangian approach for holonomic and non-holonomic constraints. The constraints applied in this paper are prescribed nodal velocities in different frames of reference, which are non-holonomic. Defining $\boldsymbol{q} = [\boldsymbol{\eta}, \boldsymbol{v}_A, \boldsymbol{\omega}_A, \boldsymbol{f}_A]$, an augmented Lagrangian is defined as

$$\mathcal{L}^*\left(\boldsymbol{q}, \dot{\boldsymbol{q}}, \right) = \mathcal{L} - k\boldsymbol{\lambda}^T\boldsymbol{\Phi} - \frac{1}{2}p\boldsymbol{\Phi}^T\boldsymbol{\Phi}, \tag{2.18}$$

with $\boldsymbol{\Phi}(\boldsymbol{q}, \dot{\boldsymbol{q}})$ the constraints imposed on the system, and $k, p \in \mathbb{R}$. These two constants have an effect on the condition number of the system of equations and the positive definiteness of the Hessian matrix, respectively. However, they do no affect the final converged result. They are thus chosen based on computational criteria. If $k = 1$ and $p = 0$, the basic Lagrange Multipliers method is recovered. This leads to a new variational expression of the action of the system:

$$\int_{t_1}^{t_2} \left[\delta\left(\mathcal{L} - k\boldsymbol{\lambda}^T\boldsymbol{\Phi} - \frac{1}{2}p\boldsymbol{\Phi}^T\boldsymbol{\Phi}\right) + \delta W\right] \mathrm{d}t = 0, \tag{2.19}$$

For an augmented Lagrangian system with non-holonomic constraints, the discrete equations of motion can be written as

$$\begin{cases} \boldsymbol{M}(\boldsymbol{q})\ddot{\boldsymbol{q}} + \boldsymbol{B}_{nh}^T(\dot{\boldsymbol{q}},t)\left(p\dot{\boldsymbol{\Phi}}_{nh}(\boldsymbol{q},\dot{\boldsymbol{q}},t) + k\dot{\boldsymbol{\lambda}}\right) &= \boldsymbol{Q} \\ k\boldsymbol{\Phi}_{nh}(\boldsymbol{q},\dot{\boldsymbol{q}},t) &= 0 \end{cases} \tag{2.20}$$

and its incremental form as

$$\begin{bmatrix} \boldsymbol{M} & 0 \\ 0 & 0 \end{bmatrix} \begin{Bmatrix} \Delta\ddot{\boldsymbol{q}} \\ \Delta\ddot{\boldsymbol{\lambda}} \end{Bmatrix} + \begin{bmatrix} \boldsymbol{C}^t & k\boldsymbol{B}_{nh}^T \\ k\boldsymbol{B}_{nh} & 0 \end{bmatrix} \begin{Bmatrix} \Delta\dot{\boldsymbol{q}} \\ \Delta\dot{\boldsymbol{\lambda}} \end{Bmatrix} + \begin{bmatrix} \boldsymbol{K}^t & 0 \\ 0 & 0 \end{bmatrix} \begin{Bmatrix} \Delta\boldsymbol{q} \\ \Delta\boldsymbol{\lambda} \end{Bmatrix} = \begin{Bmatrix} \boldsymbol{r}^* \\ -\boldsymbol{\Phi}^* \end{Bmatrix} + \mathcal{O}(\Delta^2) \tag{2.21}$$

with the following definitions:

$$\boldsymbol{r} = \boldsymbol{Q} - \boldsymbol{M}\ddot{\boldsymbol{q}} - \boldsymbol{B}_{nh}^T\left(p\boldsymbol{\Phi}_{nh} + k\dot{\boldsymbol{\lambda}}\right), \quad \boldsymbol{C}^t = -\frac{\partial \boldsymbol{Q}}{\partial \dot{\boldsymbol{q}}} + p\boldsymbol{B}_{nh}^T\boldsymbol{B}_{nh}, \quad \text{and } \boldsymbol{K}^t = -\frac{\partial \boldsymbol{r}}{\partial \boldsymbol{q}}. \tag{2.22}$$

#### 2.1.3.1 Gravity loads

Gravity loads have been embedded in the structural solver, rather than as an external force contribution. At every iteration, the "complete" mass matrix $\overline{\boldsymbol{\mathcal{M}}}$ is stored. It differs from the assembled mass matrix in the fact that the degrees of freedom corresponding to the *clamped* node in the elastic equations are not removed, and thus, the whole mass of the system is considered. Defining a vector $\boldsymbol{g}_A = [g_x|_A, g_y|_A, g_z|_A, 0, 0, 0]^T$, we can directly include the contribution of the gravity to the $\boldsymbol{\mathcal{F}}_{\text{ext}}$ vector in equation 2.12, obtained as $\overline{\boldsymbol{\mathcal{M}}}\,\boldsymbol{g}_A$. It is interesting to note that the gravity forces are the only term that introduces a dependency of the rigid body motion equations with the orientation of the $A$ frame, parametrised as $\boldsymbol{q}$.

While the converged results of this *embedded* way of including gravity loads are the same as if gravity is included as an external force contribution, the required number of elements for convergence is much lower for the embedded method. For a gravity-only loaded nonlinear beam, the number of elements to match the 10 quadratic elements solution given in Ref. [168] which is an external load implementation of gravity, is 2 for an embedded formulation.

## 2.2 Unsteady Vortex-Lattice Method

Flexible aircraft aeroelastic simulations require an integrated framework that includes flexible and rigid-body dynamics and aerodynamics. The majority of the aerodynamic solvers used in aeroelasticity based on quasi-steady approximations, and thus the wake modelling is greatly simplified and the computation times are reduced accordingly. As indicated previously, the unsteadiness of the wake can play an important role in the aeroelastic response of the aircraft. This would not be captured by quasi-steady wake approximations, specially given the characteristic low flight speeds of HALE aircraft. The UVLM is a 3-D geometrically-nonlinear aerodynamic solver [89]. It is based on potential flow theory which can accurately capture the flow unsteadiness around flexible aerodynamic surfaces at low speed. In this method, 3-D geometries are modelled as thin cambered surfaces with an attached lattice of rectilinear vortex ring elements. Each surface has a wake associated with it, idealised as a thin vortex sheet that extends a given distance —or number of panels— downstream, except for the steady case with horseshoe vortices, where an analytical extension is used and the vortices are assumed to go to infinity. Most of the initial development work for UVLM force reconstruction and its application to nonlinear aeroelasticity of flexible aircraft has been conducted by D. T. Moor [94, 201]. Geometrical nonlinearity is supported by imposing the no-penetration boundary conditions on the instantaneous deformed geometry, in contrast with the solutions based on DLM, where influence matrices are not updated. The generalisation of the steady VLM into the UVLM is achieved through the wake shedding mechanism, which is described later in section 2.2.3. Once the downwash (parametrised as circulation of the vortex rings) is determined, forces are calculated at the vertices of the aerodynamic grid and at the centre of the panels for the added mass effects, both of which will be described in section 2.2.4.

While not in the scope of this work, it is important to note that while typical UVLM problems are too expensive for some areas such as controller design, this formulation can be linearised and reduced around aerodynamically trimmed conditions when such simplification is needed [125].

### 2.2.1 Potential flow theory

Given the disturbance velocity potential $\mathbf{\Phi}$, the Laplace equation is given as

$$\nabla^2 \mathbf{\Phi} = 0. \tag{2.23}$$

The vortex ring lattice with its associated circulations is a solution to this equation, and given its linearity, every individual vortex ring is a solution itself. Assuming that a flow is incompressible, inviscid and irrotational, this solution provides a solution for the continuity equation of the flow around a lifting surface at low Mach number. The partial differential equation resulting from eq. 2.23 can be solved by placing singularities (such as vortex rings) at the boundaries. For our problem, these are lifting surfaces and the far field. The lifting surfaces require a non-penetration boundary to be imposed, that is:

$$(\nabla \boldsymbol{\Phi} + \boldsymbol{v}) \cdot \boldsymbol{n} = 0, \tag{2.24}$$

where $\boldsymbol{n}$ is the normal to the panel and $\boldsymbol{v}$ is the flow velocity at the surface, which may include incoming flow speed, atmospheric disturbances, and rigid body and elastic deformation velocity contributions. The other required boundary condition is imposed in the far field and it means that the velocity disturbance decays to 0 in the infinity:

$$\lim_{\boldsymbol{r} \to \infty} \nabla \boldsymbol{\Phi} = 0, \tag{2.25}$$

with $\boldsymbol{r}$ an arbitrary position vector with origin in the lifting body.

### 2.2.2 Discretisation and solution process

Considering a single surface without loss of generality, it is discretised in $M$ panels in the chordwise and $N$ spanwise direction. The process for generating the grid is explained later on section 2.3.1.2. The number of *bound* (lifting surface) panels and vortex rings is then $K = MN$. For a non-linearised unsteady simulation, the number of *wake* vortex rings is $K_w = M_w N$, where usually $M_w \gg M$. The full *bound* grid is described by $K_\zeta = (M + 1)(N + 1)$ lattice vertices for the surface, and $K_{\zeta w} = (M_w)(N + 1)$ for the wake, as the first set of vertices of the wake is enforced to be the located at the trailing edge vertices of the bound surface. The generalisation for multiple surfaces is straightforward. In the context of an aeroelastic solver, the aerodynamic grid is generated every FSI iteration. The UVLM formulation can be posed in any frame of reference (as long as the derivative operator is adjusted accordingly), however, the fact that the wake is naturally described in the frame of reference in which the background flow is given, makes the inertial frame of reference the simplest choice. The implementation used in this work uses a frame of reference aligned with the inertial one, but attached to the origin of the body-attached frame. The change from the custom aerodynamic frame to

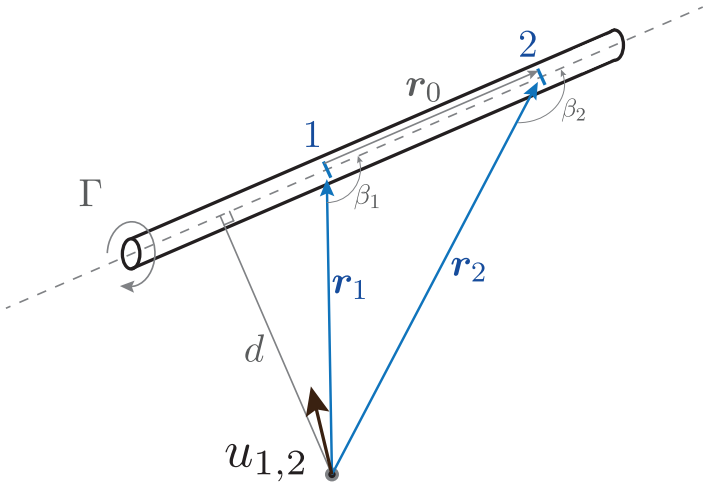



**Fig. 2.2** Definition of the variables used to evaluate the Biot-Savart law.

the inertial description is as simple as including the position of the body-attached frame of reference.

In addition to the vertex information, every vortex has a circulation, $\boldsymbol{\Gamma} \in \mathbb{R}^K$ and $\boldsymbol{\Gamma}_w \in \mathbb{R}^{K_w}$, associated with it. Once grid and circulations are defined, the instantaneous induced velocity field at a point of coordinates $\boldsymbol{x}$ can be written as

$$\boldsymbol{u}(\boldsymbol{x}, t) = \boldsymbol{\mathcal{A}}^3 \left(\boldsymbol{x}, \boldsymbol{\zeta}(t)\right) \boldsymbol{\Gamma}(t) + \boldsymbol{\mathcal{A}}^3_w \left(\boldsymbol{x}, \boldsymbol{\zeta}_w(t)\right) \boldsymbol{\Gamma}_w(t). \tag{2.26}$$

The AIC matrices $\boldsymbol{\mathcal{A}}^3$, $\boldsymbol{\mathcal{A}}^3_w$ are computed applying the Biot-Savart law. The $i$-th column of the AIC matrix is the velocity induced by the $i$-th vortex ring at the $\boldsymbol{x}_j$ spatial coordinate. The Biot-Savart law for the velocity induced by a finite linear vortex segment from $\boldsymbol{r}_1$ to $\boldsymbol{r}_2$ is

$$\boldsymbol{u}_{1,2} = \frac{\Gamma}{4\pi d}(\cos\beta_1 - \cos\beta_2)\frac{\boldsymbol{r}_1 \times \boldsymbol{r}_2}{|\boldsymbol{r}_1 \times \boldsymbol{r}_2|} = \frac{\Gamma}{4\pi}\frac{\boldsymbol{r}_1 \times \boldsymbol{r}_2}{|\boldsymbol{r}_1 \times \boldsymbol{r}_2|^2}\boldsymbol{r}_0 \left(\frac{\boldsymbol{r}_1}{r_1} - \frac{\boldsymbol{r}_2}{r_2}\right), \tag{2.27}$$

with its different terms defined in figure 2.2.

The inclusion of the unsteady wake requires to add the contributions of the previously-shedded vortex rings as an external velocity contribution, as it is independent from the current circulation values. However, the first vortex ring row of the wake, being shedded at the current time step has the same circulation value than the trailing edge bound vortex ring. Due to this, its contribution is accounted in the AIC matrix, rather than the right-hand side.

The right-hand side is then simply given as $-\boldsymbol{u}\cdot\boldsymbol{n}$ where $\boldsymbol{u}$ denotes all the external (non induced by the bound vortex rings) velocity contributions, i.e. for our case: wake-induced velocity, structural deformation velocity, rigid-body motion, atmospheric disturbance and background flow, and $\boldsymbol{n}$ is the panel normal vector. The system to solve can be written

$$\boldsymbol{\mathcal{A}}\cdot\boldsymbol{\Gamma} = \boldsymbol{u}\cdot\boldsymbol{n}. \tag{2.28}$$

### 2.2.3 Time stepping and wake shedding

The previous section gave a broad overview of the solution process. There are two main differences between the VLM and the UVLM: the time stepping and the wake shedding procedure. These two can greatly alter the results obtained and have to be considered carefully. Due to the discrete time domain formulation of the UVLM, time integration and wake shedding are intrinsically linked. The fact that the chordwise dimensions of the newly shedded wake vortex ring row is directly related to the time step adds a link between spatial and temporal discretisation. This can be written as

$$\Delta t = \frac{c}{M}\frac{1}{U_\infty}. \tag{2.29}$$

The above equation is the equivalent of CFL = 1 in a higher-fidelity finite-volume CFD context. The most important consequence of this is that a finer discretisation in the chordwise direction results in a smaller time step, which in turn results in large $M_w$ values for the same physical wake length.

Depending on the wake model fidelity, the cost of the simulation can grow very quickly, even more in coupled simulations, where the iterative nature of the strongly coupled FSI solver requires to re-evaluate the same time step a number of times. A more in-depth discussion about the consequences of this link is given in section 2.2.5.

Given a time step $\Delta t$ and an incoming velocity (excluding vortex-induced velocities) $\boldsymbol{u}$ in an unsteady simulation with no quasi-steady approximations, the wake can be modelled in a number of ways. In increasing fidelity order: 1) a fixed flat wake, 2) a moving flat wake, 3) a background-flow convected wake and lastly, 4) a full force-free wake. They are discussed next:

1. A *fixed flat wake* is the usual modelling choice in linearised aerodynamic solvers, such as the DLM. When the aerodynamic grid geometry does not vary, this usually yields good results. However, in a nonlinear aeroelasticity context, the grid deforms

and the trailing edge of the surface will not coincide with the edge of the wake. The result of this is that the Kutta condition, enforced through the overlap of the trailing edge with the wake first vortex segment, will not be imposed correctly. The results arising from this can be valid as long as structural deformations are still in the linear regime. This formulation is not expensive to evaluate, as the geometry of the wake only needs to be stored once, and thus, $\boldsymbol{\mathcal{A}}_w^3$ is constant throughout the simulation.

2. *Prescribed wake*: in sight of the unsuitability of the former wake modelling technique for the type of problems solved in this work, the wake shape can be updated to always be shedded from the trailing edge, this effectively enforcing the Kutta condition but performing a quasi-steady approximation to the convected wake kinematics. For FSI simulations where the wake induced velocity does not affect other surfaces located downstream (such as tails), this can be a suitable, relatively inexpensive approach.

3. *Convected wake*: while more suitable than the first option, the quasi-steady, updated wake modelling requires to update the $\boldsymbol{\mathcal{A}}_w^3$ term, which is expensive for the overall UVLM cost. Performing a simple extra step, the accuracy of the simulation can be greatly improved. Instead of assuming a flat wake, where the position of the vortex rings is predetermined, these can be convected with the background flow information, including atmospheric disturbances ($\boldsymbol{u}$). Two extra steps are required for this: computing the flow velocity at every $\boldsymbol{\zeta}_w$ (by interpolation or analytical expression evaluation), and convecting the vortex vertices. In this work, a simple explicit first order method is used, such that at a time step $n$:

$$\boldsymbol{\zeta}_w^{n+1} = \boldsymbol{\zeta}_w^n + \Delta t \boldsymbol{u}. \tag{2.30}$$

**All the results shown in this work have been computed using this wake convection model.**

4. *Free wake*: lastly, the wake can be convected taking into account the full velocity field, including the vortex-induced velocities. While this produces the most accurate wake dynamics, the cost of it is usually too expensive to be considered in practical cases. Furthermore, while the cost of the simulation is increased, the aerodynamic loads very rarely show any difference between the previous method and the full force-free wake model. Because the bound vortex ring circulations affect the wake

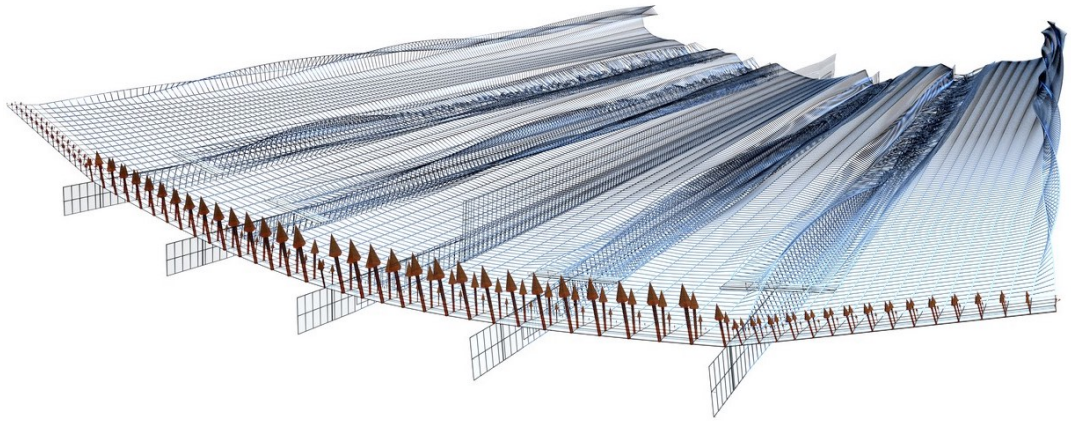


**Fig. 2.3** XHALE model (see 3.4.1) solved with a free wake.

dynamics, the most accurate way of simulating them is to perform an iterative procedure where the RHS of the UVLM is updated with the influence of the newly-position vortices. This is hugely expensive in purely aerodynamic simulations, and rarely a consideration in the case of multidisciplinary simulations. When free-wake effects are considered in this work, these are calculated using a single step and integrated using equation (2.30) with the added vortex-induced velocities in the $\boldsymbol{u}$ term. Figure 2.3 shows a free wake static aeroelastic solution.

### 2.2.4 Force calculation from UVLM solutions

Once the circulations of the vortex rings have been determined, the resulting aerodynamic forces on the surfaces have to be computed. Industry-standard aerodynamic tools in aeroelasticity focus mostly on lift prediction. However, in the case of very flexible aircraft or flapping wings, all the components of the aerodynamic force can be relevant [179] including non-stationary induced drag, which is difficult to compute in lifting-surface methods such as the UVLM due to the lack of leading-edge suction effects.

From classical potential flow theory, the leading-edge suction effect can be easily computed, as the net drag of a body inmersed in a potential flow is null —this is known as the D'Alembert's Paradox—. In unsteady problems, however, there can be a net positive or negative force in the streamwise direction. Garrick [61, 196] proposed a closed-form solution for a sinusoidally oscillating airfoil. His method is based on integrating the infinite vorticity contribution on a infinitesimal rounded leading edge.

From the Unsteady Lifting-Line Method we can extract an equivalent approach that yields a unsteady force contribution from an infinitesimal vortex segment $\delta\boldsymbol{l}$. Using the

unsteady vector form of the Joukowsky theorem, two force contributions are identified and the resulting force increment $\partial \boldsymbol{F} = \partial \boldsymbol{F}_{\text{steady}} + \partial \boldsymbol{F}_{\text{unsteady}}$. These two are given as

$$\partial \boldsymbol{F}_{\text{steady}} = \rho_\infty \Gamma (\boldsymbol{v} \times \partial \boldsymbol{l}), \tag{2.31}$$

$$\partial \boldsymbol{F}_{\text{unsteady}} = \rho_\infty \frac{\partial \Gamma}{\partial t} c (\hat{\boldsymbol{v}} \times \partial \boldsymbol{l}). \tag{2.32}$$

The direct application of the previous equations yields what we are going to call the *Joukowsky Method*, introduced in section 2.2.4.1. Another way of computing aerodynamic forces from UVLM solutions is to apply the Unsteady Bernoulli equation to the top and bottom surfaces for each panel. This yields a correct approximation for the lift, but does not account for the suction side effects that are so important for drag prediction. Including these requires some assumptions, amongst them, only accounting for the chordwise vortex segments contribution to induced drag. This is correct as long as these are perfectly aligned with the incoming flow. This method differs from the *Joukowsky* method in that it is cheaper to apply. It will be referred to as *Katz and Plotkin* method, as in Ref. [169].

#### 2.2.4.1 The Joukowsky method

The Joukowsky method relies on a direct application of equation (2.31) to the linear vortex segments that form the surface grid. This provides a *steady* force term, which then needs to be complemented with equation (2.32), which is derived from the Bernoulli equation. This last contribution will be called *unsteady* contribution and it represents the added mass effects.

For a straight vortex segment with constant circulation, as it is the case for the UVLM here explained, the steady force contribution for each of these can be given by

$$\boldsymbol{f}_{\text{steady}} = \rho_\infty \left( \tilde{\boldsymbol{v}}_s \times \boldsymbol{\ell}_s \right) \Gamma_s \tag{2.33}$$

where the $s$ subindex denotes that specific segment, $\boldsymbol{\ell}$ is the segment vector and $\tilde{\boldsymbol{v}}$ denotes the total relative segment-fluid velocity at its centre point (including vortex-induced velocities).

Obtaining the velocity at the centre point of the segment usually requires an interpolation step. This implementation of the UVLM relies on a vertex-based storage, where grid geometry, relative velocities and output forces are stored at the vertices. Due to the relatively small size of the grid panels, a simple first order interpolation scheme is used

to transfer the velocities from the two extrema of the segment to its centre. Once the force at the centre of the segment is calculated, it is then lumped into the grid corner points belonging to the segment.

Obtaining $\tilde{\boldsymbol{v}}$ requires to evaluate the vortex-induced velocities on the centre of the segment. It can be shown that induced velocities from colinear vortex-segments are null, which settles one of the potential problems. However, the self-induced velocities of a vortex segment are infinite. This is approached by modifying the Biot-Savart law with a numerical epsilon or a vortex core cut-off.

The unsteady force contribution $\boldsymbol{f}_{\text{unsteady}}$ (equation (2.32)) is obtained from the application of Bernoulli's equation accross the surface grid. This contribution is calculated at the $k$-th collocation point as

$$\boldsymbol{f}_{\text{unsteady}_k} = \rho_\infty A_k \hat{\boldsymbol{n}}_k \dot{\Gamma}_k, \tag{2.34}$$

where $A$ denotes the panel area and $\dot{\Gamma}$ is the time derivative of the vortex circulation with respect to time. This quantity needs to be numerically evaluated, usually with backwards finite differences schemes. The implementation used in this work uses a first order method. In aerodynamic-only simulations, this term does not have a detrimental effect on the stability of the numerics. However, FSI simulations are more sensitive to it, as the circulation $\Gamma$ is a function of the structural displacements and velocities, and a numerical derivative approximation can be noisy or not smooth enough. This can result in numerical divergence in the FSI iterative scheme. Section 2.2.5.3 provides a description of the numerical methods used in order to improve the behaviour of the simulation in simulations where the added mass effects are relevant.

#### 2.2.4.2 The generalised Katz and Plotkin method

This generalised Katz and Plotkin method is a simplified version of the previous one. While its implementation is more complex, some of the quantities required for it have been already computed during the solution process, and the fact that a modified velocity field is used for computing the forces allows to avoid evaluating the vortex core routine at —or next to— its singularity.

This section borrows the notation used in [89]. The local lift force contribution for the $ij$-th panel, applied at its collocation point is given by

$$L_{ij}^{\text{local}} = \rho_\infty \left[ \left( \boldsymbol{U}_{ij}^u + \boldsymbol{U}_{ij}^w \right) \cdot \boldsymbol{\tau}_{ij}^c \frac{\Gamma_{ij} - \Gamma_{i-1,j}}{\Delta c_{ij}} + \left( \boldsymbol{U}_{ij}^u + \boldsymbol{U}_{ij}^w \right) \cdot \boldsymbol{\tau}_{ij}^s \frac{\Gamma_{ij} - \Gamma_{i,j-1}}{\Delta b_{ij}} + \frac{\partial \Gamma_{ij}}{\partial t} \right] A_{ij} \cos \alpha_{ij}, \tag{2.35}$$

where $\boldsymbol{U}_{ij}^u$ is the relative velocity between fluid and lattice collocation point due to aerodynamic inputs, that is, $\boldsymbol{U}_{ij}^u = \boldsymbol{U}_{ij}^{\text{ext}} - \dot{\boldsymbol{\zeta}}_{ij}^{\text{col}}$, where the first term of the right hand side is the external fluid velocity, and the second term is the motion of the collocation point. $\boldsymbol{U}_{ij}^w$ is the velocity contribution from the wake vorticity at the collocation points. The panel tangential vectors in the chordwise and spanwise directions are noted as $\boldsymbol{\tau}^c, \boldsymbol{\tau}^s$, while the local angle of attach is given by $\alpha_{ij} = \tan^{-1}\left( \boldsymbol{U}_{ij}^u \cdot \hat{\boldsymbol{n}}_{ij} / \boldsymbol{U}_{ij}^u \cdot \hat{\boldsymbol{\tau}}_{ij}^c \right)$.

In this method, the induced drag is computed by considering only the contribution of the chordwise vortex segments. This greatly reduces the cost of the computation, but it is not an accurate representation of the potentially complex lateral motion and cross-wind situations that can arise from a low-speed, low-altitude flight. With an orthogonal projection operator that projects the panel normal onto the plane perpendicular to the relative velocity [62] $P_{\hat{\boldsymbol{U}}_{ij}^u} = \mathbb{I}_3 - \hat{\boldsymbol{U}}_{ij}^u \hat{\boldsymbol{U}}_{ij}^{u^T}$, we can define the local drag term as:

$$D_{ij}^{\text{local}} = \rho_\infty \left[ - \left( \boldsymbol{U}_{ij}^{bc} + \boldsymbol{U}_{ij}^w \right) \cdot \left( P_{\hat{\boldsymbol{U}}_{ij}^u} \cdot \hat{\boldsymbol{n}}_{ij} \right) \left( \Gamma_{ij} - \Gamma_{i-1,j} \right) \Delta b_{ij} + \frac{\partial \Gamma_{ij}}{\partial t} A_{ij} \sin \alpha_{ij} \right]. \tag{2.36}$$

Where the *bc* superscript denotes the chordwise-oriented vorticity contribution to the induced velocity on the surface only. The total aerodynamic force contribution from each panel is $\boldsymbol{f}_{ij} = D_{ij}^{\text{local}} \hat{\boldsymbol{U}}_{ij}^u + L_{ij}^{\text{local}} P_{\hat{\boldsymbol{U}}_{ij}^u} \hat{\boldsymbol{n}}_{ij}$. This can then be interpolated into the grid vertices.

## 2.2.5 Practical considerations

### 2.2.5.1 Biot-Savart law desingularisation

The Biot-Savart law, included in equation (2.27), yields the velocity induced by a vortex segment (straight in this simplified version) on a point not belonging to the vortex segment. Already in section 2.2.4.1 it has been pointed that the application of the UVLM in general, and specifically the Joukowsky force calculation method can result in incorrect results when evaluating the induced velocities in neighbouring points of the vortex segment. The vortex core model is important also when considering surface-wake

interactions. If one of the vortex segments that belongs to the wake is close enough to the collocation point or mid-point of a segment belonging to a lifting surface, the induced velocity estimation can yield incorrect results. It is also worth noting that if the wake is simulated as a force free wake, the vortex rings can deform, rotate and translate to the point where vortex segment intersections can happen, this causing more numerical issues in the wake convection if the vortex core is not correctly modelled.

This can be solved in multiple ways [152]. In this work, two approaches are explored: a vortex core cutoff and a regularisation term in the Biot-Savart evaluation.

The vortex cutoff method is the most used in usual UVLM evaluations. It offers a number of benefits. Amongst them: simplicity of implementation and intuitivity of the method. Given a value of the vortex radius cutoff, $r_c$, the cutoff method can be summarised as

$$\boldsymbol{u}_{1,2}\big|_{\text{cutoff}} = \begin{cases} 0 & \text{if } d \leq r_c \\ \boldsymbol{u}_{1,2} & \text{if } d > r_c \end{cases}, \tag{2.37}$$

where $\boldsymbol{u}_{1,2}$ is the induced velocity given by equation 2.27 and $d$ is the distance between the point and the vortex segment given by

$$d = \frac{\left\| \left(\boldsymbol{r}_p - \boldsymbol{r}_1\right) \times \left(\boldsymbol{r}_p - \boldsymbol{r}_2\right) \right\|}{\left\| \left(\boldsymbol{r}_2 - \boldsymbol{r}_1\right) \right\|}, \tag{2.38}$$

with its terms defined in figure 2.4.

The vortex cutoff presents a number of drawbacks. The most important are: the $r_c$ value case (and spatial discretisation) dependency and the added cost of computing the distance and adding a conditional clause in the Biot-Savart law evaluation. The author has explored several possibilities in order to reduce the case and discretisation dependency of the vortex cutoff value. The most reliable has been to, when generating the aerodynamic grid, keep track of the shortest vortex segment of the lattice. For simple simulations, a $r_c$ value smaller than 1/4 of the length of the shortest is usually a good starting value, as the collocation point is located at $3c/4$ of the panel downstream. However, this value only guarantees correct results when no wake vortex ring intersects any other vortex ring. Cases when this can happen are: free-wake simulations —especially in cases with large wing tip vortices or large flight speed variations—, and aircraft with lifting surfaces located downstream of another —such a conventional T-tail configurations—.

Concerning the computational cost of computing the distance between the collocation point and the vortex segment, if done carefully, only a small overhead is added. However,

the addition of an *if* statement and a varying return point depending on it makes the optimisation of the routine more challenging. It is worth noting that for average simulations, the Biot-Savart law is called millions of times, so a small gain in performance cannot be dismissed, as it will have a much greater effect in the overall code. On one side, the computation of the distance and the inclusion of the *if* clause, together with the different return statement, increase the execution time, while the fact that the induced velocity does not need to be computed in the cases where $d \leq r_c$ reduces it with respect to a pure Biot-Savart evaluation. While highly dependant on the compiler and the code quality, the overhead is usually greater than the savings. This is due to the fact that the distance and the extra *if* need to be computed every single iteration, while the savings only exist in the less common case $d \leq r_c$.

The second option to improve the robustness of the Biot-Savart evaluation is the addition of an epsilon value $\epsilon_c$ in the denominator of the terms of equation (2.27) that present a singularity when $d \to 0$ as explained in [156]. The reformulated Biot-Savart evaluation results:

$$\boldsymbol{u}_{1,2} = \frac{\Gamma}{4\pi} \frac{\boldsymbol{r}_1 \times \boldsymbol{r}_2}{|\boldsymbol{r}_1 \times \boldsymbol{r}_2|^2 + \epsilon_c} \boldsymbol{r}_0 \left( \frac{\boldsymbol{r}_1}{r_1 + \epsilon_c} - \frac{\boldsymbol{r}_2}{r_2 + \epsilon_c} \right). \tag{2.39}$$

This model requires $\epsilon_c$ to be defined, which can be more difficult than the former approach due to its lack of physical meaning. However, tests performed for this work show that the value is more independent from the case being run than for the $r_c$ case. A numerical investigation has shown that values of $\epsilon_c$ are usually larger than $10^{-10}$ and small enough to not affect the solution. A typical value would be $\epsilon_c = 10^{-6}$. While a custom choice of $r_c$ can make a simulation more robust, the epsilon method requires less user input and most of the times the results are comparable to the ones obtained with the former.

Figure 2.4 shows how the induced velocity in the $d \to 0$ region is modified for the original Biot-Savart and the two modified versions.

#### 2.2.5.2 Spatial and temporal discretisation considerations

For a typical strongly coupled scheme, as implemented in this work, if the constraint in equation (2.29) is not complied with, the accuracy of the method is compromised[2].

[2]Later work by A. Muñoz-Simón included in SHARPy has circumvented this potential issue, which is critical for the simulation of High Aspect-ratio Wind Turbines. The modified scheme has not been used here.

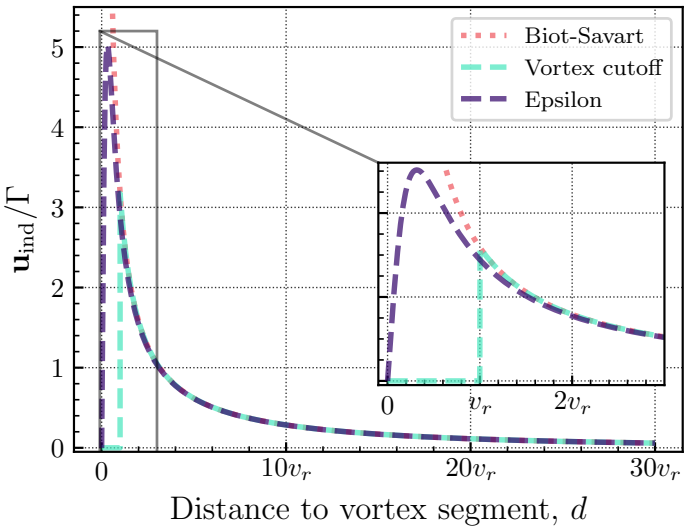



**Fig. 2.4** Biot-Savart law including the different core models implemented for this work.

Although some small deviation can be worth considering for some cases in order to reduce the cost of the simulation, most of the cases will benefit from a problem discretisation that agrees with the constraint. There are two cases where fully complying with it is not possible or practical: different chord length surfaces and varying incoming velocity due to rigid-body motions or non-constant background flow as a function of time or span position.

In the first case, the most obvious solution is to select the number of chordwise panels $M$ of the surfaces such that the panel chordwise lengths ($\Delta c$) are similar (considering uniformly-distributed panels). Precision requirements might require a minimum number of chordwise panels in order to accurately capture the camber shape, especially in the case of non-symmetric airfoils with complex cambers, such as reflex airfoils, or control surfaces. When this is the case, this minimum number of chordwise panels sets a constraint in $\Delta c$ for the rest of the panels. If the chord lengths are substantially different between surfaces, the resulting number of panels in the larger ones can be too large to be simulated with the available resources. Usually the best solution is to set a compromise between strongly enforcing equation (2.29), accuracy requirements and the available resources.

When a case with large relative flow speed variations with time is being simulated, equation (2.29) cannot be enforced throughout the whole simulation without modifying the time step in runtime. Due to the coupled nature of the simulations here presented, it has been deemed too cumbersome to contemplate the implementation of runtime-varying time steps. Instead, the time step is selected so that it is during a baseline flight condition

that has been identified as a common flight condition with important aerodynamic effects. For example, in a catapult take off simulation, the time step is chosen such that equation (2.29) is enforced at cruise velocity, which will be close to the maximum speed (and maximum dynamic pressure) seen in the case. In the case where the incoming flow speed is a function of time and spanwise position, such that the flight of a large span vehicle in a 3-D turbulent field, this choice is more difficult still. As long as the velocity differences are not excessive, this will not cause noticeable problems for very flexible aircraft simulations. However, other fields where SHARPy is applied, such as wing engineering can be more dependent on this, and other coupling schemes are necessary.

#### 2.2.5.3 Added mass filtering for coupled simulations

As mentioned in Section 2.2.4.1, the unsteady force contribution given by equation (2.32) is proportional to the time derivative of the vortex ring circulations. In aerodynamics-only simulations, this term does not affect the numerical properties of the method in any significant way. However, once this is integrated in a coupled framework, the addition of a force term that is directly related to a numerically-evaluated derivative and which, in turn, affects the structural deformations and deformation velocities, can result in divergent behaviour even in simple cases [8].

In sight of the typical physical properties of HALE aircraft –low structural mass, high aspect ratio, potentially large unsteady structural deformations–, the added mass effects can be an important factor in their dynamic response. For cases where the added mass effects are not relevant enough to be worth considering –featuring only slow inputs such as atmospheric disturbances and control surface deflections–, this contribution can be omitted. This results in faster convergence of the FSI solver, which translates in important computational cost savings.

For the other cases featuring potentially strong added mass effects, this term needs to be included. If the unsteady force term is directly implemented as described in equation (2.32), most of the cases show a numerically-induced high-frequency vibration directly after the first time step, or after a short simulation time. The solution to these two problems needs to be approached separately but simultaneously, once the first time step instability has been solved, it is likely that it will still diverge artificially during the simulation process if the second problem is not taken care of.

The divergence after the initial time steps is due to a potentially large step in loading between the static and the first dynamic solution. This can be due to an inaccurate trim of the aircraft, different modelling techniques between static and dynamic solvers

–most usually due to the wake treatment–, or different velocity field in both fields. This last is typical of open-loop simulations in continuous turbulence, where the aircraft is trimmed for average conditions at a certain altitude, but once it is immersed in the velocity field, the trim conditions are not exact. It is beneficial to introduce a certain number of time steps where the added mass effects are ignored until the effect of the step in loading is damped. A typical number of time steps with no added mass would be around 2 chords of flow convection. The second way of improving the convergence of the coupled solver is to introduce the added mass effects in a ramp during a number of FSI iterations every time step. Ramping the unsteady forces has a strong effect in the number of FSI iterations needed for convergence, as the solution cannot be considered converged until as many FSI iterations as steps in the ramp have been run. However, this method notably improves the numerical behaviour of the coupled solver during the whole simulation. The recommendation is to use this only in cases with strong added mass effects that show a tendency to numerically diverge.

In most of non-trivial cases studied in this work, the added mass term cannot be dismissed, and the two solutions explained previously do not completely eliminate the possibility of the added mass-induced instability of the method. This has been solved by filtering the $\dot{\Gamma}$ estimation obtained through finite-differences. The filter used in this work is a Wiener filter [101] as implemented in SciPy. It is used to estimate and filter the noise from signals, and the most interesting of its properties for our problem is that the noise can be estimated without any input from the user, given a long enough signal. The required input is a window width for the smoothing step, called here $\eta$, which has to be odd and $\geq 3$. At every sample point $i \in [0, \cdots, N-1]$ of the discrete signal $w$, the local mean $\mu_i$, and variance $\sigma_i$, are computed in their neighbourhood:

$$\mu = \frac{1}{\eta} \sum_{j \in [-\eta/2, \cdots, \eta/2]} w_{i+j}, \text{ and } \sigma^2 = \frac{1}{\eta} \sum_{j \in [-\eta/2, \cdots, \eta/2]} w_{i+j}^2 - \mu_i^2. \tag{2.40}$$

The estimate returned by the filter is given as:

$$\bar{w}_i = \mu_i + \frac{\sigma_i^2 - \nu^2}{\sigma_i^2} \left( w_i - \mu \right). \tag{2.41}$$

In our case, where the noise variance $\nu^2$ is not known, it is approximated as the average of all the local estimated variances.

In all the cases shown in this work, the window width has been set at 7 time steps (around 2 convected chord lengths) with the results not showing a strong dependency

with similar values of the parameter. In order to accurately estimate the signal noise, the raw, unfiltered $\dot{\Gamma}$ history is fed into the algorithm, including the current time step to be filtered. After running the filter, the last element of the resulting time history is used to compute the unsteady component of the aerodynamic forces.

## 2.3 A strongly coupled FSI algorithm

Once the structural and aerodynamic solvers are described, coded and validated, a coupling scheme is required. While weakly coupled FSI schemes are simpler to implement, the higher accuracy for the same time discretisation from the strongly coupled schemes results in smaller aerodynamic problems to be solved, and thus, overall faster results for the same accuracy, despite the overhead of multiple structural and aerodynamic evaluations. Based on the available literature and in-house experience, a strongly coupled formulation has been chosen for this work, while keeping in mind that a weakly coupled simulation can be run in a strongly coupled solver with a simple setting modification. This work presents strongly coupled results except in a few cases where it will be explicitly noted.

The coupling algorithm is based on a Block Gauss-Seidel iteration scheme. The main advantage of this is the ease of coupling of the structural and aerodynamic solvers, which can be treated as black boxes as long as the force and mapping procedure between them is defined. Figure 2.5 shows a high-level flowchart for the coupling algorithm implementation. The following sections describe the most relevant blocks.

### 2.3.1 Force and displacements mapping

A fundamental part of every coupled solver is the force and displacement mapping procedure. The implementation here described relies on a coincident spanwise discretisation, where the lattice cross-sections are defined for each node in the finite element description of the structure, as shown in figure 2.6.

#### 2.3.1.1 Force mapping

Let's assume that the aerodynamic forces and moments vectors at the $i, j$ vertex for the grid, $\boldsymbol{f}_a^{(ij)}$ and $\boldsymbol{m}_a^{(ij)}$, have already been lumped at the vertices of the aerodynamic grid $\boldsymbol{\zeta}$. With the coincident mesh assumption, no interpolation step is necessary. In typical FSI solvers where the aerodynamic discretisation is generally not coincident with the

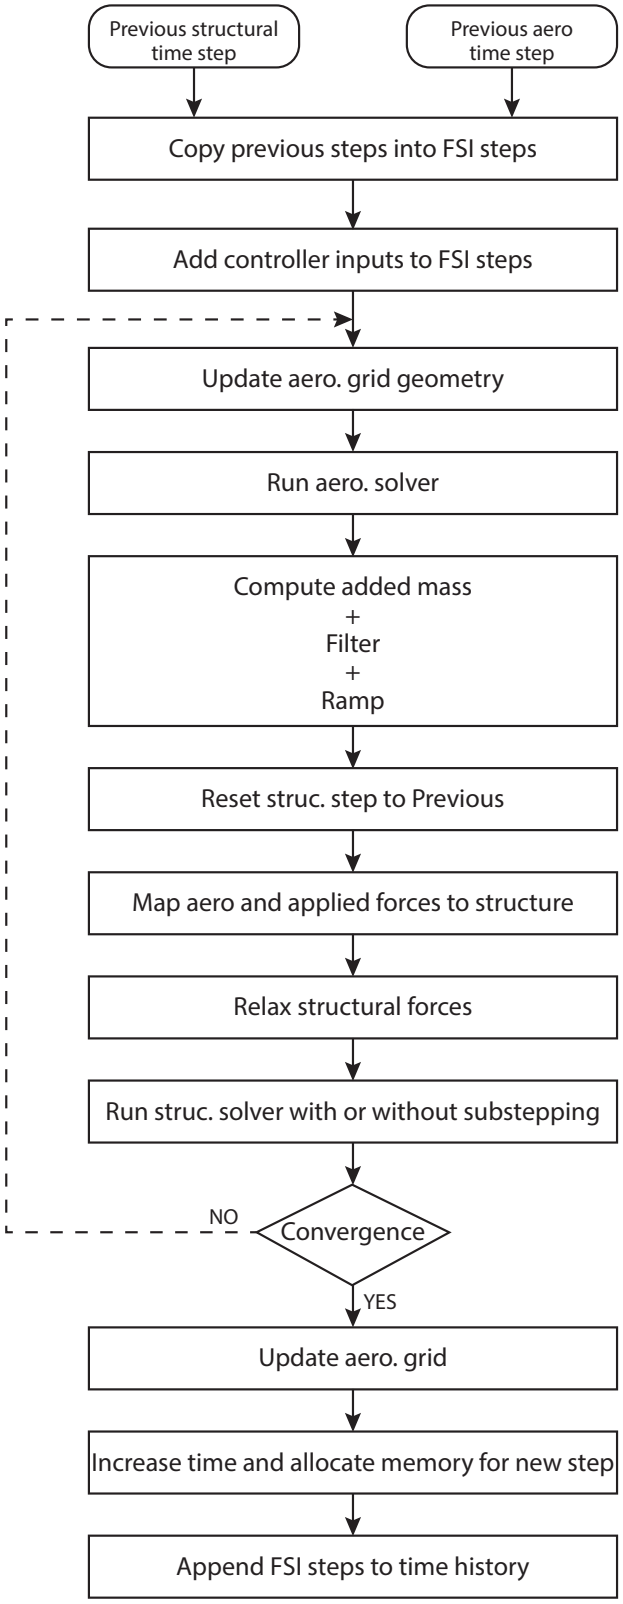



**Fig. 2.5** Flowchart of the BGS strongly coupled scheme described in this work.

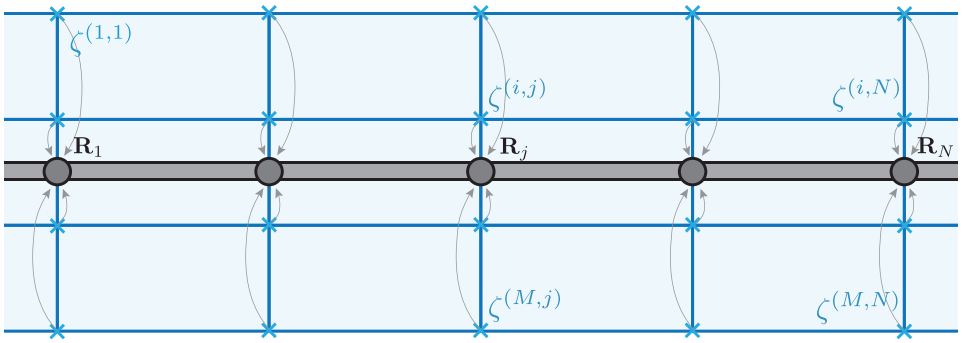



**Fig. 2.6** Force mapping procedure for coincident structural ($\boldsymbol{R}_j$) and aerodynamic ($\boldsymbol{\zeta}^{(i,j)}$) meshes.

structural one, transferring the aerodynamic loads is not a trivial exercise. A good review of the requirements and some common FSI force interpolation methods is given in Ref. [54].

The structural applied loads are implemented as a follower force. As such, the most suitable frame of reference in which to describe them is in material, $B$, frame of reference. This is in contrast with the aerodynamic solver formulation, which provides the aerodynamic loads in inertial frame. The resulting structural nodal forces and moments at a structural node $j$, $\boldsymbol{f}_s^{(j)}$ $\boldsymbol{m}_s^{(j)}$ are given for a single surface by

$$\boldsymbol{f}_{sB}^{(j)} = \sum_{i=0}^{M} \boldsymbol{C}^{BA}\left(\boldsymbol{\Psi}_j\right) \boldsymbol{C}^{AG}\left(\boldsymbol{q}\right) \boldsymbol{f}_{aG}^{(ij)} \tag{2.42}$$

$$\boldsymbol{m}_{sB}^{(j)} = \sum_{i=0}^{M} \boldsymbol{C}^{BA}\left(\boldsymbol{\Psi}_j\right) \boldsymbol{C}^{AG}\left(\boldsymbol{q}\right) \left( \boldsymbol{m}_{aG}^{(ij)} + \left( \boldsymbol{\zeta}_G^{(ij)} - \boldsymbol{C}^{AG^T}\left(\boldsymbol{q}\right) \boldsymbol{R}_A^{(j)} \right) \times \boldsymbol{f}_{aG}^{(ij)} \right). \tag{2.43}$$

It is important to note that these equations are suitable for structural nodes belonging to a single lifting surface. However, it is a common situation to have a shared node between a number of surfaces, e.g. a T-tail model, where the node at the tip of the fin is also the root of the horizontal stabilisers. When this is the case, the aerodynamic forces are simply lumped into the same node, paying special attention to updating the rotation matrix $\boldsymbol{C}^{BA}$ for every surface, as they will belong to a different beam element, and thus, a run through the master-slave tree of the structure is needed to extract the different CRV for each lifting-surface associated node.

A master-slave structure is necessary when dealing with 1-D beam formulations where overlapping nodes are not supported, like in the case of this GECB solver. A 1-D

structure is defined in terms of nodes an elements. More specifically, a quadratic element is defined by 3 nodes, and in general, only the central one is not shared with other elements, the exception to this being free tips and some bifurcations. When generating the data structure containing the FEM information, every time a node is shared between two or more elements, one of them is chosen to be the *master*. As such, the rotation information used for constructing the submatrices of the FEM model related to the node comes from its master. The slave element information is used to reconstruct the structural state in the postprocessing stage.

#### 2.3.1.2 Displacement and velocities mapping and grid generation

Thanks to the simple aerodynamic grid required by the UVLM, it is easier and quicker to build the grid at every time step instead of deforming it as usual for CFD-based FSI simulations [32]. Under the hypothesis of rigid cross-sections, the grid generation step is reduced to a multiple 3-D rotations exercise. In addition to the structural positions and rotations, the grid generation supports additional inputs such as twist, $\gamma$, and sweep, $\lambda$ angles. These are given with respect to the $B$ frame of reference components —twist is implemented as a rotation of the camber line with respect to the $x$ axis, while the sweep angle is a rotation around the local $z$ axis of the material frame of reference, as defined in Fig. 2.1—.

In conventional UVLM solvers using the Generalised Katz and Plotkin force calculation method (see section 2.2.4.2), it is important that the grid is aligned with the incoming flow in order to obtain a relatively accurate induced drag estimation. However, this is much more difficult once the incoming flow is not uniform or presents large sideslip angles. Using the Joukowsky force calculation method does not require to have the grid aligned with the incoming flow. This in turn improves the robustness of the grid generation process.

In addition to updating the grid vertices coordinates, the structural deformation linear and angular velocities need to be mapped into the aerodynamic grid, as well as the control surface deflection velocities. We split the velocity in three components as:

$$\dot{\boldsymbol{\zeta}}_{ij} = \dot{\boldsymbol{\zeta}}_{ij}\Big|_R + \dot{\boldsymbol{\zeta}}_{ij}\Big|_\Psi + \dot{\boldsymbol{\zeta}}_{ij}\Big|_{cs}. \tag{2.44}$$

First, the linear structural deformation velocities are converted and mapped from the structural nodes onto their corresponding grid vertices:

$$\dot{\boldsymbol{\zeta}}_{ij}\Big|_R = \boldsymbol{C}^{GA}\left(\boldsymbol{\Psi}\right)\dot{\boldsymbol{R}}_A. \tag{2.45}$$

The rotational velocities need to be transformed from the time derivative of the CRV description to a conventional $\boldsymbol{\omega}$ one. From Géradina and Cardona [62]:

$$\boldsymbol{\omega} = \boldsymbol{T}^T\left(\boldsymbol{\Psi}\right)\cdot\dot{\boldsymbol{\Psi}}, \tag{2.46}$$

where $\boldsymbol{T}$ is the tangential operator, defined in equation (2.6). The contribution of the rotational deformations to the grid velocities is then given by

$$\dot{\boldsymbol{\zeta}}_{ij}\Big|_\Psi = \boldsymbol{C}^{GA}\left(\left(\boldsymbol{T}^T\left(\boldsymbol{\Psi}\right)\cdot\dot{\boldsymbol{\Psi}}\right)\times\left(\boldsymbol{\zeta}_{ij} - \boldsymbol{C}^{AG^T}\left(\boldsymbol{q}\right)\boldsymbol{R}_k^A\right)\right). \tag{2.47}$$

The last component of the aerodynamic grid velocities is the control surface deflection velocity, $\dot{\boldsymbol{\zeta}}\Big|_{cs}$. To avoid redundant inputs, the control surface angular deflection velocity is obtained by first order finite differences. For smooth control surface deflection histories, higher order FD do not provide an evident benefit, although more complicated inputs have not been considered.

### 2.3.2 Relaxation factor of the FSI solutions

In coupled simulations where iterative procedures are performed, most of the times it is suitable to include a relaxation factor in the solution from one solver before introducing it into the other. In this work, we introduce a relaxation factor $\omega \in [0, 1)$. It is then applied at a time step $t$ to the aerodynamic forces and moments vector $\boldsymbol{f}_{\text{aero}}^t$:

$$\boldsymbol{f}_{\text{aero}}^t\Big|_{\text{relax}} = (1-\omega)\ \boldsymbol{f}_{\text{aero}}^t\Big|_{\text{original}} + \omega\ \boldsymbol{f}_{\text{aero}}^{t-1}\Big|_{\text{original}}. \tag{2.48}$$

The last term $\boldsymbol{f}_{\text{aero}}^{t-1}\Big|_{\text{original}}$ can be substituted by $\boldsymbol{f}_{\text{aero}}^{t-1}\Big|_{\text{relax}}$ in order to avoid storing extra time steps in memory, but the solution will convergerce slightly slower.

Most of the typical coupled simulations will benefit from a relaxation factor of around 0.3, while the most challenging ones —usually close to the flight envelope edge— can require values of up to 0.8.

### 2.3.3 Structural solver substepping

Section 2.2.5.2 already made reference to the link that the UVLM introduces between the spatial and temporal discretisations. Strongly coupled FSI schemes require the structural and aerodynamic solvers to be run using the same time step. The deciding factor in number of chordwise panels and time step is usually the accuracy of the lift and drag predictions, which in turn can have a even larger effect once the structural deformations are accounted for. However, high frequency structural dynamics can be necessary to converge the structural solver, even when they will not affect the aerodynamics. This is specially important in the case of transient response analysis. The structural integrator, a Newmark-$\beta$ method is not unconditionally stable when there is a high-frequency component in the response. That is why a numerical damping parameter is added, affecting mostly high-frequency modes.

These cases with higher-frequency content require a small $\Delta t$ in the structural solver, but it is not necessary, nor practical, to increase $M$ in the aerodynamic discretisation accordingly. The resulting size of the aerodynamic problem would make it impractical. A solution to this problem that does not require to run a weakly coupled simulation based on structural solver substepping integrated in a strongly coupled FSI algorithm has been developed and implemented. A simple plot providing an outline of it can be found in Figure 2.7.

In this procedure, the structural solver time discretisation is refined, such that: $N_{\text{sub}}\Delta t_s = \Delta t_a$. Structural sub-steps are run in the FSI iteration scheme, while aerodynamic loads, applied forces and time-dependent constraints on the structural solver are linearly interpolated during FSI iterations between the previous time step ($Q_{s-1}$) and the last FSI iteration of the current aerodynamic time step ($Q_s^k$). Applied and aerodynamic forces are introduced as follower forces and their orientation is updated every structural sub-step. Flight dynamics variables are integrated every structural time step, increasing their accuracy.

This is only implemented in the FSI algorithm, which means that the structural solver does not need to be modified at all to be used in this kind of simulations. Every FSI iteration includes then $N_{\text{sub}}$ structural steps and an aerodynamic one. The rest of the FSI solver is left untouched, including the convergence check. The result of this is a more accurate structural response, specially in the higher-frequency region of the spectrum, while adding only a slight overhead, as the most expensive part of the simulation is still the UVLM. Using this prevents the use of weakly coupled formulations with lower

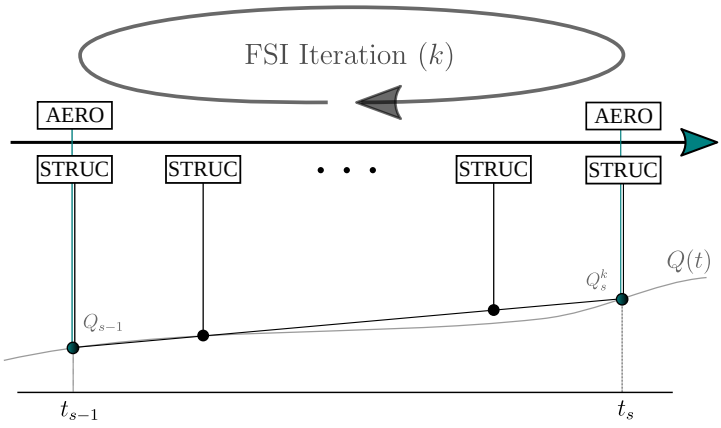



**Fig. 2.7** Subiteration process in the strongly coupled FSI scheme.

accuracy, avoids the increase in size of the UVLM problem, and improves the accuracy with which multibody constraints are imposed on the problem.

Implementing this method simply requires expanding the block corresponding to the structural solver in Figure 2.5 with the one presented in Figure 2.8.

### 2.3.4 FSI convergence check

The definition of a robust convergence assessment method is not a trivial exercise. The potential range of variation in the typical values between various cases is of several orders of magnitude.

The convergence criterium used in this work is strict enough to capture small discrepancies between iterations, but at the same time some checks have been implemented to avoid the stalling of the solution. The basic error definition in this work is defined as:

$$\epsilon(\bullet^k) = \frac{\left\|\bullet^k - \bullet^{k-1}\right\|}{\left\|\bullet^0\right\|}, \tag{2.49}$$

where $\bullet^k$ denotes a certain vector quantity at an iterative step $k > 0 \in \mathbb{N}$. This definition is used to compute the error associated to the structural state vector and its derivative, $\epsilon([\boldsymbol{\eta}\ \boldsymbol{\beta}]^T)$ and $\epsilon([\dot{\boldsymbol{\eta}}\ \dot{\boldsymbol{\beta}}]^T)$. Incorporating the deformation velocity into the convergence threshold sets a much more strict criterium, as derivatives tend to fluctuate faster than the original values, specially when using a method such as the Newmark-$\beta$. In some cases, the threshold normalisation might result in large final values, which in turn causes large number of iterations or a stall in the convergence. In order to avoid this, the absolute residual $\epsilon^+$ is defined as the non normalised version of the one in eq. (2.49). The values

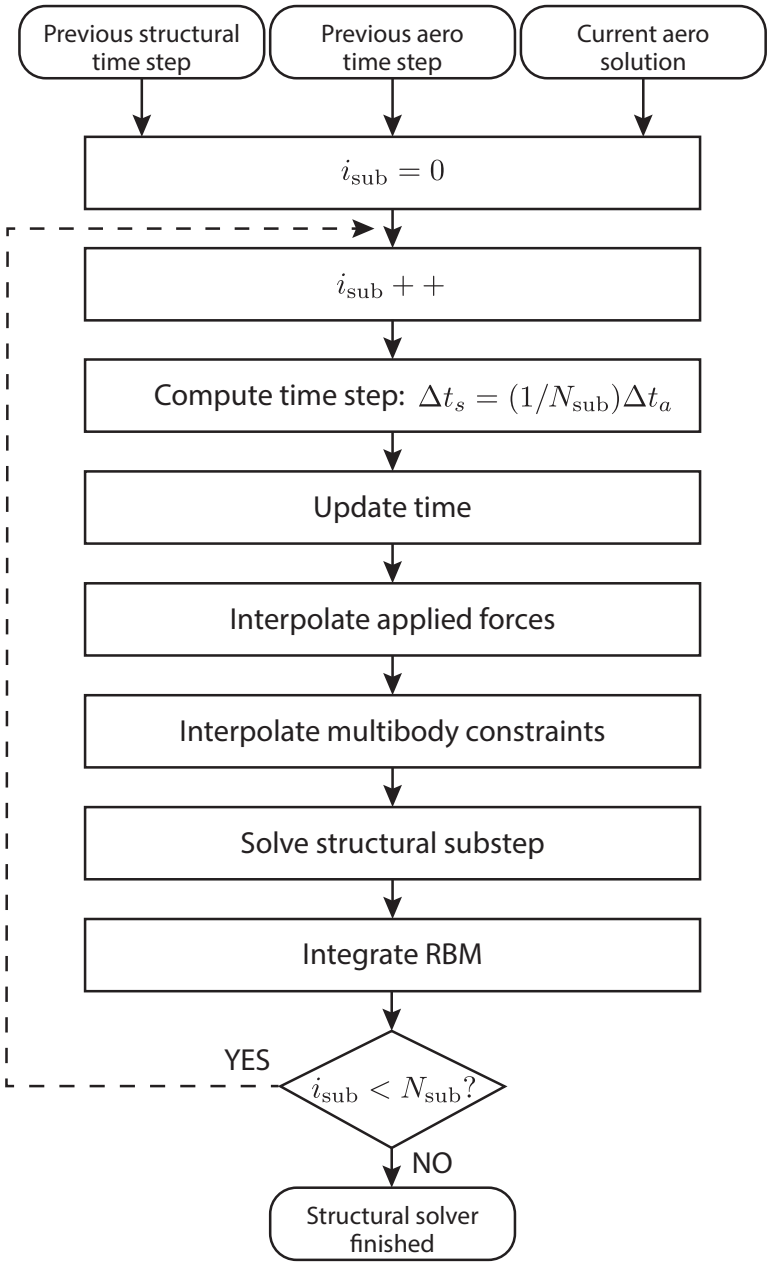



**Fig. 2.8** Flowchart for the custom strongly coupled FSI substepping procedure.

resulting from the application of $\epsilon^+$ are also checked against the user-given convergence threshold.

Finally, a minimum number of FSI iterations is required. If the added mass contribution ramping is being used (see Section 2.2.5.3), the minimum FSI iterations needs to be equal or greater than the ramp length.

### 2.3.5 Longitudinal trim solution

An important effect of aeroelasticity in general is the change in trim condition of the aircraft due to the elasticity of the structure. This is even more relevant in cases with large structural deformations even while in cruise flight. Typical trim algorithms rely on the aerodynamic and control derivatives of the aircraft found through asymptotic analysis. While this is a totally valid option for up to moderately flexible aircraft, it might pose some problems when dealing with very flexible aircraft that can show non monotonic trends in their derivatives, such as control reversal. A potential solution for this is to update the linear model every few iterations of the trim solver. However, it can be quite costly to generate a fully coupled linear model for a one-off evaluation. The approach proposed in this work is a simplified Newton-Raphson method with first order finite differences gradient evaluation, where the Jacobian is assumed to be predominantly diagonal.

For a longitudinal trim problem, we assume the inputs to be: angle of attack, $\alpha$, elevator deflection, $\delta$ and thrust per motor, $T$. The magnitudes being monitored are: resultant vertical and horizontal force in the longitudinal axis and pitching moment in inertial frame of reference. Computing the whole Jacobian matrix, which would be composed of 9 terms, 6 of them independent and requiring at least two evaluations each, would result in 12 static aeroelastic solutions per iteration of the solver. Instead, the Jacobian matrix is assumed to be diagonal, i.e. the vertical force is only affected by angle of attack, and so on. This is not a bad approximation in the longitudinal trim case, as the coupling between variables is a second order effect at least, and reduces the required number of function calls to 4 per Jacobian evaluation: one baseline and one extra per input variable.

In addition to the diagonal Jacobian matrix, instead of the deflection angle of the elevator, the solver is formulated in terms of its incidence angle, defined as: $\delta_i = \alpha + \delta$. This is done internally, and the input and output is given in terms of tail deflection. The effect of this is to detach the angle of attack from the pitching moment as much

as possible. This internal change of variables greatly accelerates the convergence of the trim solver, as the assumption of the diagonal shape of the Jacobian is reinforced.

# Chapter 3

# Aeroelastic solver verification studies

This chapter introduces the numerical methods used in this work. Their theoretical basis was already described in doctoral theses previously conducted at the Loads Control and Aeroelastics lab at Imperial College London [123, 69, 170]. While the current implementation draws from their work, the solvers have been rewritten in order to improve their robustness, efficiency and increase the size of the problems that can be tackled. This work has been released under the name SHARPy (Simulation of High Aspect-Ratio planes in Python), and can be freely downloaded under an Open Source license. SHARPy features a versatile user interface, where complex models can be described. The implementation is designed to be highly modular, accepting custom functionality that can be added with no need to modify the base code. The amount of rewritten code, changes and new functionality introduced in the existing methods justifies the need for a verification exercise of all the components of the code.

First, some generic results that are typical in the literature are shown in order to verify the implementation of the independent solvers used in this work. This includes geometrically nonlinear beam validation cases, VLM comparisons with the TORNADO [114] code, and static aeroelasticity based on a simple wing defined in [173].

The second set of results is derived from a collaborative work between the author and the team at University of Michigan. The target of such work was to consolidate a description of the X-HALE, create a model based on it for SHARPy and UM/NAST (Nonlinear Aeroelastic Simulation Toolbox) [180], compare both codes in a static setting first, and then move onto dynamic simulations of "1-cos" gust encounters. The X-HALE is a nonlinear aeroelasticity testbed for experimentation and validation developed at University of Michigan [84]. It has 5 tails arranged along the span of a very high aspect-ratio wing. It features several characteristics that are representative of state of the art

| **Property** | **Value** |
|---|---|
| Length | 5 [m] |
| $EA$ | $4.8 \cdot 10^8$ N |
| $GA$ | $3.231 \cdot 10^8$ N |
| $GJ$ | $1.0 \cdot 10^6$ Nm |
| $EI$ | $9.346 \cdot 10^6$ Nm |
| Tip mass | 61162.08 kg |

**Table 3.1** Geradin and Cardona [62] beam properties used for dead loads validation in a nonlinear context.

| Model | $\Delta R_1$ (m) | $\Delta R_3$ (m) | $\Psi_2$ (rad) |
|---|---|---|---|
| Simpson 2 3-noded elements | 0.550 | -2.070 | -0.6576 |
| 2 3-noded elements | 0.596 | -2.159 | -0.6718 |
| Simpson 5 3-noded elements | 0.589 | -2.144 | -0.6700 |
| 5 3-noded elements | 0.598 | -2.160 | -0.6720 |
| Simpson 10 3-noded elements | 0.596 | -2.159 | -0.6719 |
| 10 3-noded elements | 0.597 | -2.160 | -0.6720 |
| Geradin and Cardona | | -2.159 | -0.6720 |

**Table 3.2** Current solver results, Simpson [168], and Geradin and Cardona [62] results comparison.

HALE configurations, such as distributed mass and propulsion, high flexibility —specially in torsion—, and low controllability, this being even more challenging as the aircraft can experience unstable flight dynamics from a lateral mode. A full description of the X-HALE and the discretisation used in the comparison exercise is given in section 3.4.1.

Last, a simple representative HALE model is described. This will be used in the next chapters.

## 3.1 Geometrically-Exact Beam validation

### 3.1.1 Nonlinear structures: dead loads

The case here replicated for static solution of gravity (dead) loads, and nonlinear deformation validations is the one presented in Geradin and Cardona [62]. Being the gravity loads implementation a new implementation (see section 2.1.3.1), it has been deemed necessary to validate the results.

This case is a simple isotropic cantilever beam with a mass in its tip. The beam properties are given in Table 3.1.

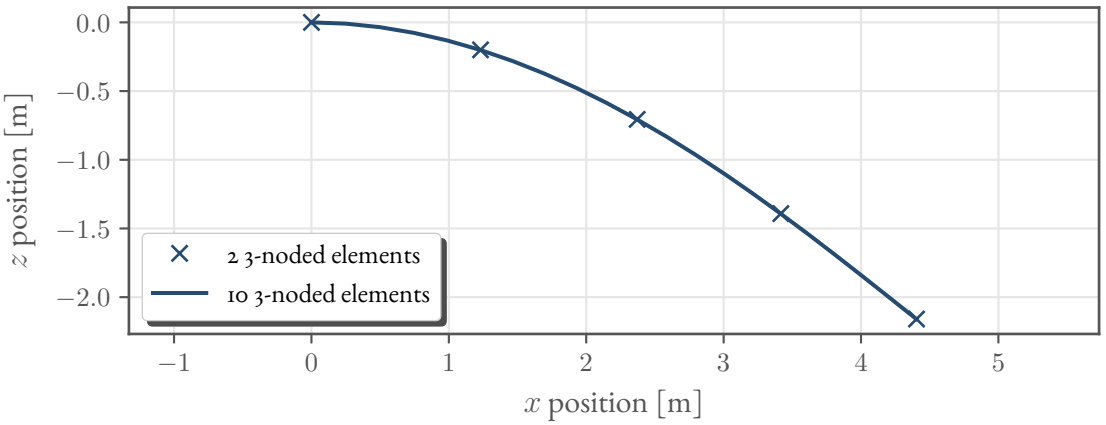



**Fig. 3.1** Comparison of beam shape with 2 and 10 3-noded elements for the Geradin and Cardona beam.

From Table 3.2 we can see how the new implementation requires even fewer elements than the previous iteration of the library. At the same time, the agreement with the literature is excellent, even with only 2 quadratic elements (Fig. 3.1).

### 3.1.2 Nonlinear follower loads

In this implementation of the structural solver, the follower loads are handled differently from the gravity loads. Due to this, a different validation procedure for the follower loads has been considered necessary. Here, we are going to use the same beam properties from the previous validation case (Table 3.1), although no gravity is included. The applied load is going to be a moment in the normal direction at the wing tip. The movement is thus planar. The angle of the tip element of a cantilever beam under a normal moment applied in this same tip node has an analytical solution [91] and is given by

$$\Psi_2 = \frac{LM_2}{EI}. \tag{3.1}$$

In order to validate the code, an increasing follower moment is going to be applied on the tip of the beam, up until the value required for having $\Psi_2 = 2\pi$ in 3 increments. The beam will be discretised using 5 or 20 3-noded elements. The values of $\Psi_2$ are $\mathcal{O}(10^{-15})$ rad for the fully loaded 5 and 20 elements cases. This is effectively machine precision when working with double precision variables. The results for 5 and 20 3-noded elements are shown in figure 3.2.

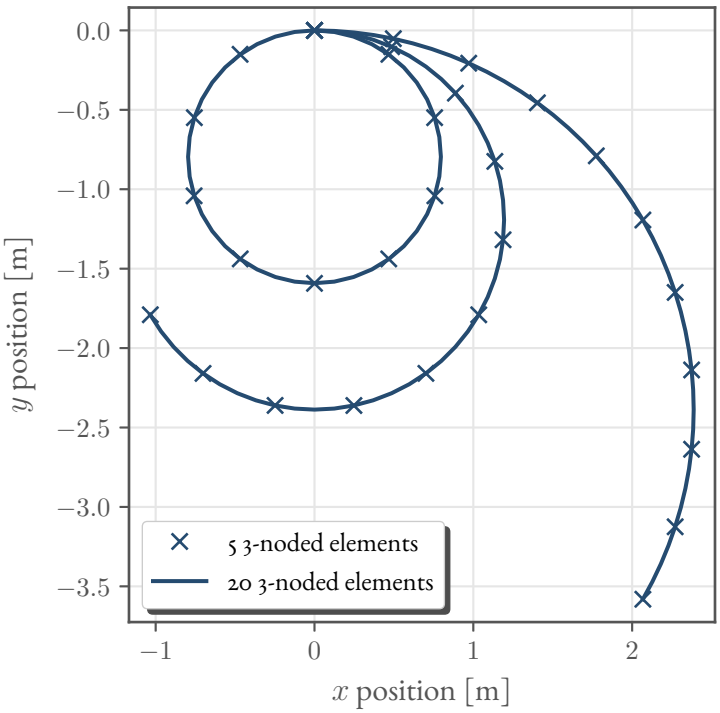



**Fig. 3.2** Beam shape for follower moments, including 5 and 20 3-noded elements results.

### 3.1.3 Rigid body rotation objectivity errors

The interpolation of rotations in the FEM discretisation should enforce its objectivity, that is, the invariance of the strain field under arbitrary rigid-body rotations. Table 3.3 shows the obtained results for the beam used in section 3.1.1 when loaded with a 3000 kN follower tip force. The beam deforms in the $x - y$ plane and the force is applied in the $z$ direction. The resulting tip element CRV is then rotated back to the 0° angle position in order to be able to compare results.

Table 3.3 shows how a rotation of the beam with respect to the inertial frame of reference introduces some numerical error. This error increases with larger angles and coarser discretisations. Simpson [168] also compares the 3-noded elements with the simpler 2-noded ones and shows that the improvement in numerical accuracy makes the 3-noded elements the best option for large deformations and large rigid-body motions models. The order of magnitude of the errors in $\Psi_1$ and $\Psi_3$ are the same for the current work and for Simpson.

| Angle [°] | Model | $\Psi_1$ [rad] | $\Psi_2$ [rad] | $\Psi_3$ [rad] |
|---|---|---|---|---|
| 0 | 10 3-noded elements | $0.00 \times 10^{0}$ | $-2.76$ | $0.00 \times 10^{0}$ |
| 90 | 10 3-noded elements | $3.29 \times 10^{-5}$ | $-2.76$ | $-3.42 \times 10^{-4}$ |
| 180 | 10 3-noded elements | $-4.40 \times 10^{-5}$ | $-2.76$ | $-4.62 \times 10^{-4}$ |
| 0 | 50 3-noded elements | $0.00 \times 10^{0}$ | $-2.76$ | $0.00 \times 10^{0}$ |
| 90 | 50 3-noded elements | $8.45 \times 10^{-17}$ | $-2.76$ | $-4.22 \times 10^{-7}$ |
| 180 | 50 3-noded elements | $-7.19 \times 10^{-8}$ | $-2.76$ | $-6.67 \times 10^{-7}$ |

**Table 3.3** CRV values for a beam under a follower tip force of 3000 kN with different azimut angle between the $G$ and $A$ frames of reference.

## 3.2 Static aerodynamic loads

### 3.2.1 3-D effects validation

In order to assess the accuracy of the aerodynamic model, a comparison with data from Katz & Plotkin [89] has been performed.

The discretisation size has been chosen based on [123]. $M = 8$ panels are used in the chordwise direction, while the spanwise discretisation has $N = 10 \times \mathcal{AR}$. Finer discretisations do not change noticeably the results. Considering this, and for consistency's sake, the discretisaton has been kept the same as in Murua's work.

The data shown in figure 3.3 features the slope of the $C_L - \alpha$ curve for different aspect ratios ($x$ axis) and sweep angles ($\Lambda$).

#### 3.2.1.1 Planar wing validation against TORNADO

Tornado [113] is a vortex lattice implementation focused on linear aerodynamics. It features a fixed-wake VLM, which is essentially the same theory in the current solver for static solutions.

As a test case, the HALE wing described by Smith et al. in [175] is used. Using its underformed shape, the discretisation has been made finer and finer until the memory requirements for Tornado were too high for being run in the workstation. The flight conditions and basic geometry parameters are given in table 3.4. A geometry output and a circulation distribution contour map is shown in figure 3.4.

The results are shown in figure 3.5. It can be seen how both drag and lift converge to the same values for Tornado and the UVLM presented here. However, it is important to note several points. First, the drag estimate is not going to be an accurate prediction, because panel methods only predict induced drag, and do not take into account any

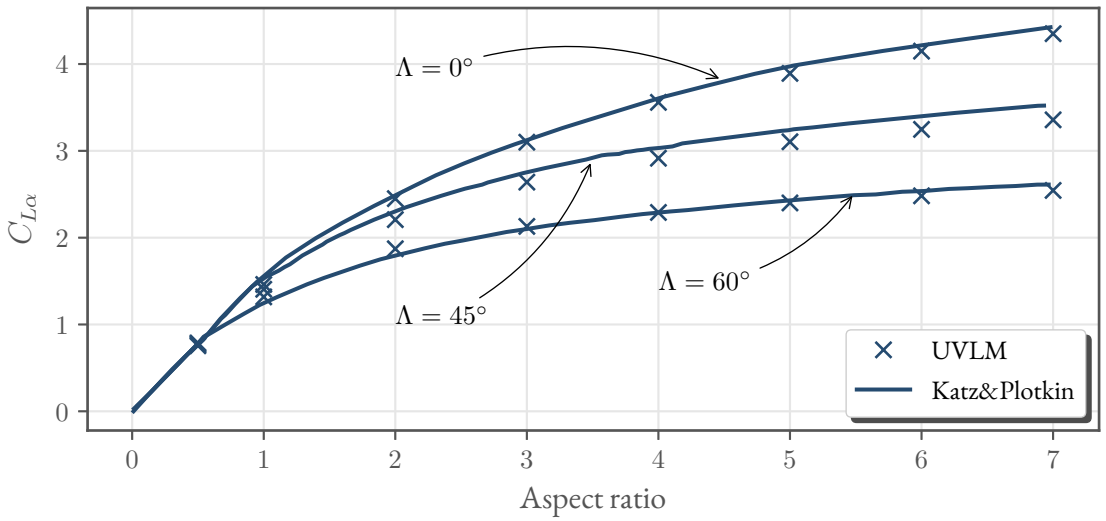



**Fig. 3.3** Comparison of obtained $C_{L\alpha}$ vs. aspect ratio for a planar wing between the current UVLM implementation and Katz & Plotkin [89]

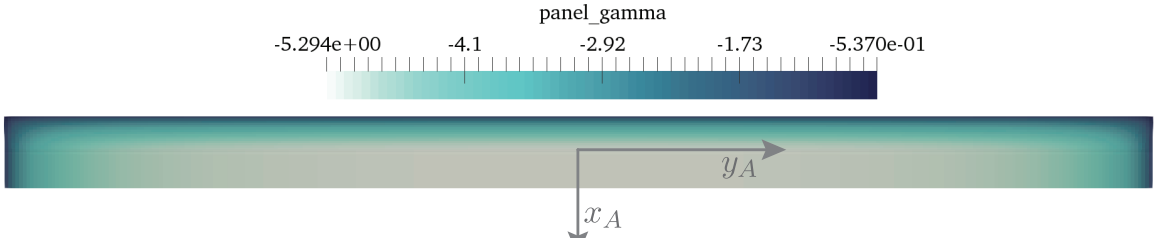



**Fig. 3.4** Contour map of panel circulation for the planar wing defined in table 3.4.

| | | |
|---|---|---|
| semi-span $b/2$ | 16 | [m] |
| sweep $\Lambda$ | 0 | [deg] |
| freestream speed $U_\infty$ | 25 | [m/s] |
| air density $\rho$ | 1.225 | [kg/m$^3$] |
| angle of attack $\alpha$ | 4 | [deg] |
| sideslip $\beta$ | 0 | [deg] |

**Table 3.4** Input data for the Tornado comparison case.

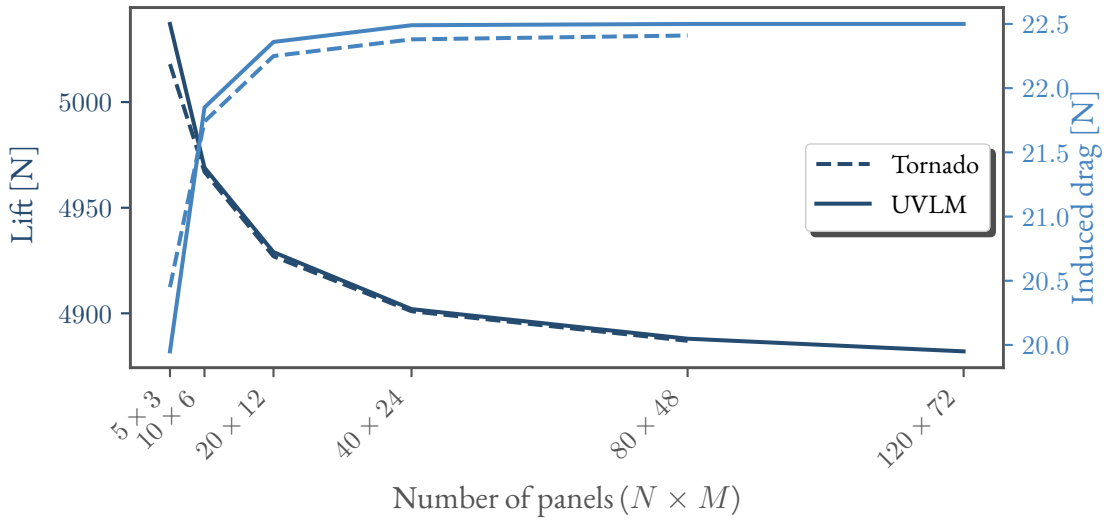



**Fig. 3.5** Comparison of results between the current UVLM solver and Tornado for the HALE wing [175].

other drag sources. Second, the memory requirements of Tornado were 10 times higher than what this UVLM required. For example, the $160 \times 96$ case could not be run because Tornado needed to allocate a 150 Gb matrix, while this solver only needed 15 Gb. This is even more remarkable if we take into account that no image method is available in our UVLM implementation. So, effectively, the 15 Gb required were for a case double the size than the one simulated in Tornado. A possible reason for this is the implementation of the steady wake using horseshoe vortex elements used in our solution, which noticeably reduces the computation time and memory requirements.

## 3.3 Nonlinear static aeroelastic results

This last test validates the coupling strategy and the force/displacement mapping. It consists on a flexible high-aspect ratio wing (see figure 3.6) derived from featured in section 3.6 and already used in the former subsection.

The flight conditions are indicated in table 3.5. Data from Smith [175] is used in the no gravity case, while the data from Simpson [168] provides the comparison data for the gravity case. It is important to note that the gravity implementation of Simpson has not been validated in aeroelastic cases nor it is clear that a convergence study was conducted for the discretisation.

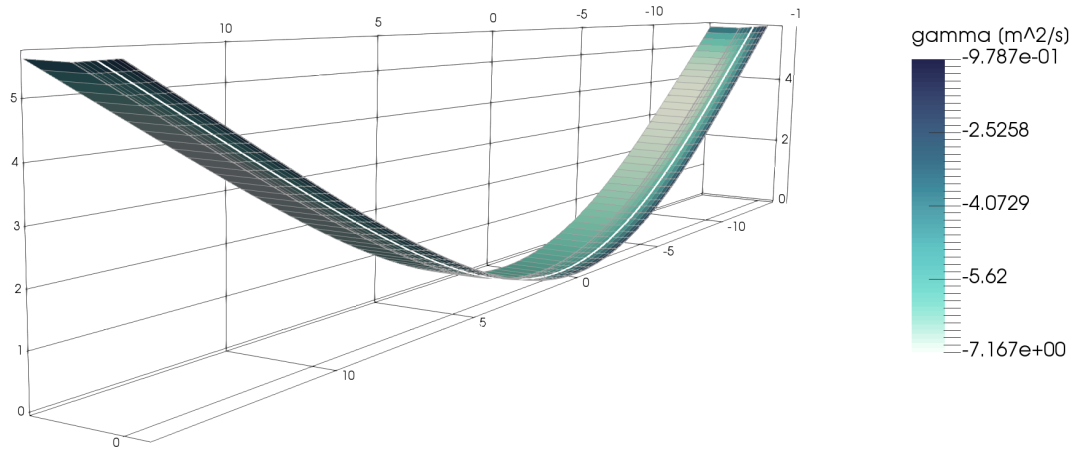



**Fig. 3.6** Aeroelastically trimmed configuration with panel circulation contours for the Smith wing, $\alpha = 4°, g = 0$

| | | |
|---|---|---|
| Freestream speed $U_\infty$ | 25 | m/s |
| Air density $\rho$ | 0.08891 | kg/m$^3$ |
| Angle of attack $\alpha$ | 2 or 4 | deg |
| Sideslip $\beta$ | 0 | deg |

**Table 3.5** Flight conditions for the flexible HALE wing case.

Zero-gravity case results are presented in figure 3.7. The $\alpha = 4°$ case shows almost overlapping lines, while the case $\alpha = 2°$ presents a small difference near the wing tip. Smith's dataset is based on a potential method, while they also show another set of results using a finite volumes inviscid CFD. For that dataset, the tip deflection is even larger than the one predicted by SHARPy. We can then infer that our potential flow aerodynamics implementation is slightly more accurate than the one by Smith. In addition to the beam deflection plots, figure 3.6 presents a typical SHARPy output. In this case, it is the $\alpha = 4°$, no-gravity case.

The gravity case has been simulated using a gravity acceleration value of $g = 9.754\,\mathrm{m\,s^{-2}}$, which is the one corresponding to the altitude for the given air density value.

## 3.4 X-HALE aeroelasticity

Static simulations have been performed with the X-HALE model in order to validate the fundamentals of the aeroelastic solver in a more complex setup. First, the X-HALE

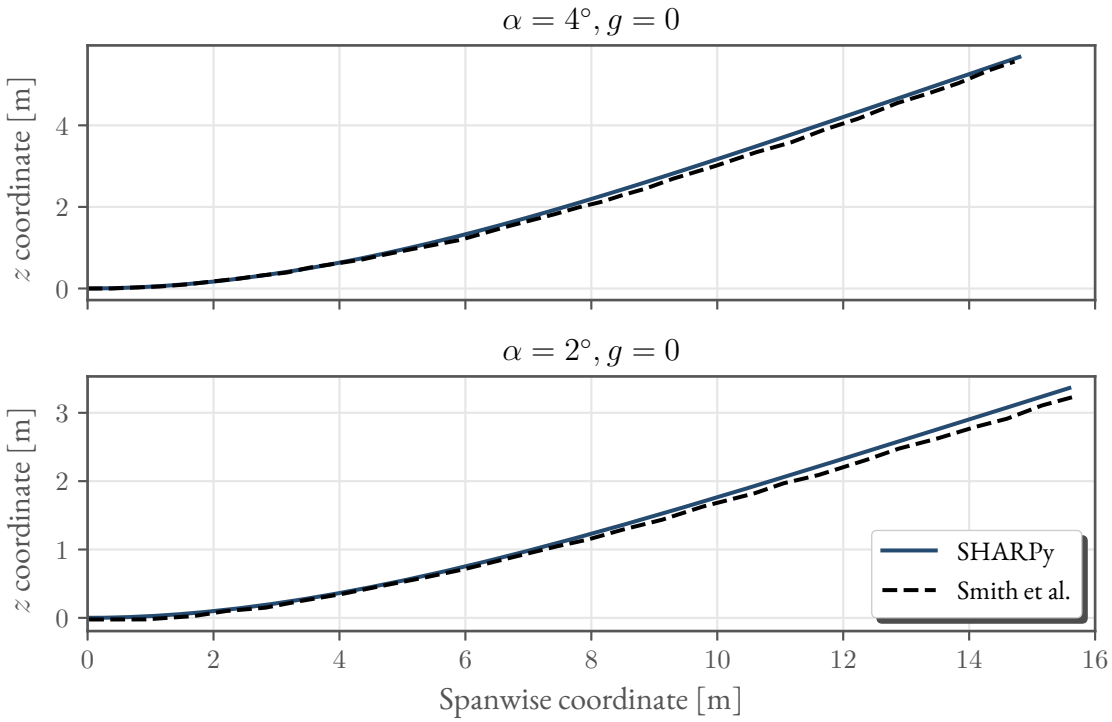



**Fig. 3.7** Vertical wing deflection ($g = 0$) comparison between Smith et al. [175], and SHARPy

model is fully described. Then, the standalone structural and aerodynamic solutions are independently compared, before the full aeroelastic trim equilibrium of the aircraft is sought.

### 3.4.1 The X-HALE model

The X-HALE UAS (Fig. 3.9) has been developed at the Active Aeroelasticity and Structures Research Laboratory at the University of Michigan. It serves as a nonlinear aeroelastic testbed intended for code validation and controller development for very flexible aircraft. Its main wing is composed of 6 panels of constant structural and aerodynamic properties, with a 10 degree dihedral region on the outer panels. The centre tail is actuated along the longitudinal axis and it can be rotated to act either as a horizontal or vertical tail. It is known that flipping the centre tail modifies the stability properties of the aircraft: when the tail is in its horizontal position, the aircraft presents an unstable leaf mode. The remaining four tails (two on each semi-span) act as elevators, deflecting uniformly. Three ventral fins, between the wing and the center

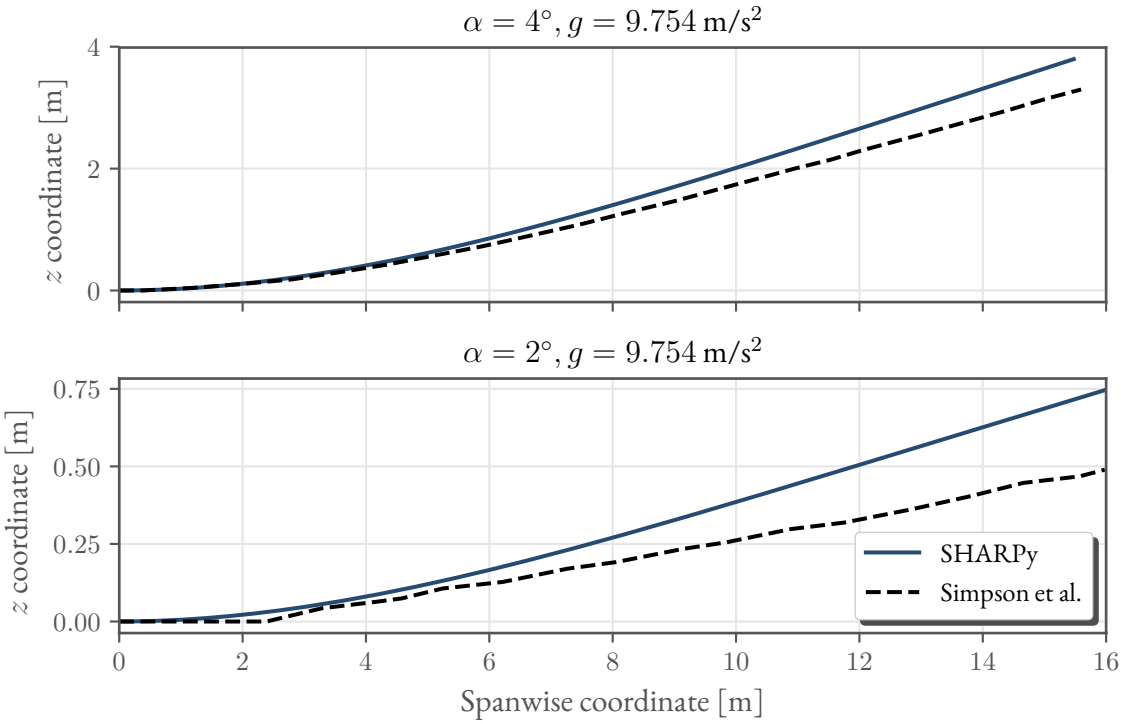



**Fig. 3.8** Vertical wing deflection ($g$ =9.754 m s$^{-2}$) comparison between Simpson et al. [168], and SHARPy

and inboard tails, are also present to augment the lateral stability of the X-HALE. Five propellers, one located at each pod, provide thrust. These pods also contain the avionics equipment, battery packs, instrumentation, as well as the landing gear. X-HALE presents an aspect ratio of 30, with a relative tip deflection of approx. 18% of its semi-span under nominal cruise conditions [84] and even larger tip deflections during manoeuvres or gust conditions.

The X-HALE features a number of challenging features for aeroelastic modeling and simulation. Firstly, it has a relatively complex geometry with multiple intersecting lifting surfaces. Secondly, all its primary subcomponents are very lightweight and their flexibility need to be considered, which is done here with beam elements with elastic and inertia coupling. Thirdly, as wings are span-loaded, it has a rather involved mass distribution.

The first bending mode of the aircraft has a natural frequency close to 0.60 Hz, which results in a strong coupling between the elastic and the flight dynamic modes. The relatively low stiffness in torsion of the main wing presents another set of challenges. The influence of the properties of the tails and their trim angle is more remarkable than for



**Fig. 3.9** The X-HALE on the ground (left) and in flight shortly after take off (right). Images from [85].

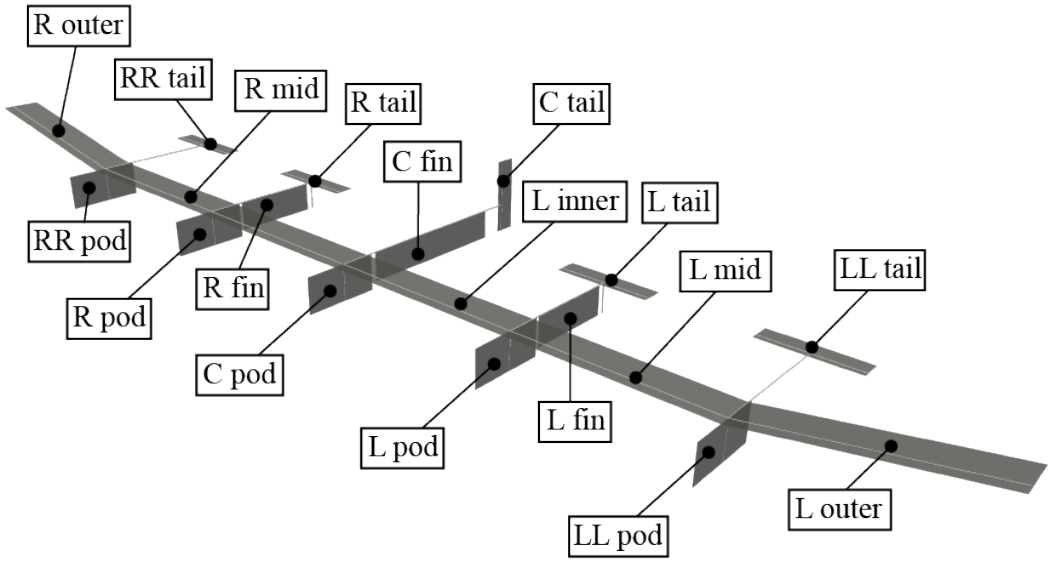



**Fig. 3.10** Nomenclature followed in the X-HALE model description.

standard configurations, as they also affect the lift distribution of the wing by modifying its twist. It is also worth noting that being the tail in the wake of the main wing, the wake modeling is of relevance, specially in dynamic simulations [126].

The stiffness distribution of the X-HALE model is given in Table 3.6 and based on the nomenclature indicated in Figure 3.10. For briefness sake, only right wing information is given, but the whole model is easy to reconstruct considering that the stiffness and distributed mass data in Table 3.7 are symmetrical. It is important to note that in that table, the mass for the pods is set at $10^{-8}$ kg m$^{-1}$. However, this is set this way because the whole mass contribution of the pods comes from the lumped (discrete) masses. The lumped masses for every pod are given in table 3.8, next, the aerodynamic

| Section [-] | EA [Nm$^2$] | GJ [Nm$^2$] | EI$_y$ [Nm$^2$] | EI$_z$ [Nm$^2$] | k$_{14}$ (axial-IP) [Nm$^2$] | k$_{34}$ (OOP-IP) [Nm$^2$] |
|---|---|---|---|---|---|---|
| R inner | $2.14\times10^6$ | $5.93\times10^1$ | $1.12\times10^2$ | $6.35\times10^3$ | 0 | $4.63\times10^1$ |
| R mid | $2.14\times10^6$ | $5.93\times10^1$ | $1.12\times10^2$ | $6.35\times10^3$ | 0 | $4.63\times10^1$ |
| R outer | $2.14\times10^6$ | $5.93\times10^1$ | $1.12\times10^2$ | $6.35\times10^3$ | 0 | $4.63\times10^1$ |
| fuselage | $5.39\times10^7$ | $5.39\times10^7$ | $5.39\times10^7$ | $5.39\times10^7$ | 0 | 0 |
| tail | $3.21\times10^6$ | $2.14\times10^1$ | $9.10\times10^1$ | $4.27\times10^3$ | $7.44\times10^4$ | $2.26\times10^-6$ |
| fin | $5.39\times10^7$ | $5.39\times10^7$ | $5.39\times10^7$ | $5.39\times10^7$ | 0 | 0 |

**Table 3.6** Stiffness data for X-HALE model.

| Section [-] | Mass [kg/m] | $I_{xx}$ [kgm] | $I_{yy}$ [kgm] | $I_{zz}$ [kgm] | $I_{yz}$ [kgm] | $x_{cg}$ [m] | $y_{cg}$ [m] | $z_{cg}$ [m] |
|---|---|---|---|---|---|---|---|---|
| R inner | $3.94\times10^{-1}$ | $8.1\times10^{-4}$ | $1.22\times10^{-5}$ | $7.97\times10^{-4}$ | $6.5\times10^{-6}$ | 0 | $2.94\times10^{-2}$ | 0 |
| R mid | $3.94\times10^{-1}$ | $8.1\times10^{-4}$ | $1.22\times10^{-5}$ | $7.97\times10^{-4}$ | $6.5\times10^{-6}$ | 0 | $2.94\times10^{-2}$ | 0 |
| R outer | $5.0\times10^{-1}$ | $8.1\times10^{-4}$ | $1.22\times10^{-5}$ | $7.97\times10^{-4}$ | $6.5\times10^{-6}$ | 0 | $2.14\times10^{-2}$ | 0 |
| fuselage | 0.0429 | $2.91\times10^{-9}$ | $1.46\times10^{-9}$ | $1.46\times10^{-9}$ | 0 | 0 | 0 | 0 |
| tail | 0.2614 | $1.6\times10^{-4}$ | $2.910\times10^{-6}$ | $1.57\times10^{-4}$ | 0 | 0 | 0.0144 | 0 |
| pod | $1\times10^{-8}$ | $2\times10^{-8}$ | $1\times10^{-8}$ | $1\times10^{-8}$ | 0 | 0 | 0 | 0 |
| C fin | 0.5092 | $3.19\times10^{-3}$ | $9.34\times10^{-5}$ | $3.28\times10^{-3}$ | 0 | 0 | 0 | 0 |
| L/R fin | 0.3208 | $8.17\times10^{-4}$ | $5.88\times10^{-5}$ | $8.76\times10^{-4}$ | 0 | 0 | 0 | 0 |

**Table 3.7** Distributed mass of the X-HALE model.

model geometry is given in table 3.9, and its discretization is shown in table 3.10. Lastly, for the results shown in this work, the wake length is set at 4 m, independently of the chordwise discretization.

### 3.4.2 Structural cases

In order to compare the solutions provided by the structural solver, three sets of simulations have been conducted: 1) gravity forces-only loading, 2) a vertical follower force, and 3) an out-of-plane bending follower moment at the wing tip. They are all applied on half a X-HALE vehicle clamped at the wing midpoint. These cases have been chosen to validate several aspects of the solver: the geometrical nonlinearity treatment, the follower forces and moment implementation, and the gravity and stiffness distributions of the model.

Figure 3.11 shows the wing spar in-plane and out-of-plane displacements in a gravity-only loading case. For the sake of clarity only the wing spar position is shown, even though the case includes the whole aircraft structural model.

Next, the half-aircraft model is subject to two load cases, first, a wing tip vertical follower force ($F_z = 15\,\mathrm{N}$) and second, a wing tip out-of-plane bending moment ($M_y =$

| Mass [kg] | $x_{cg}$ [m] | $y_{cg}$ [m] | $z_{cg}$ [m] | $I_{xx}$ [kgm$^2$] | $I_{xy}$ [kgm$^2$] | $I_{xz}$ [kgm$^2$] | $I_{yy}$ [kgm$^2$] | $I_{yz}$ [kgm$^2$] | $I_{zz}$ [kgm$^2$] |
|---|---|---|---|---|---|---|---|---|---|
| **C pod** | | | | | | | | | |
| 0.3746 | 0 | 0.1 | 0 | $1.15\times10^{-3}$ | 0 | 0 | $8.90\times10^{-4}$ | 0 | $8.90\times10^{-4}$ |
| 1.0462 | $3.97\times10^{-3}$ | 0.0612 | -0.0168 | $1.48\times10^{-2}$ | $2.32\times10^{-4}$ | $2.27\times10^{-5}$ | $2.82\times10^{-3}$ | $4.50\times10^{-4}$ | $2.50\times10^{-4}$ |
| 0.023 | 0 | 0.260 | -0.023 | 0 | 0 | 0 | 0 | 0 | 0 |
| **L pod** | | | | | | | | | |
| 0.548 | -0.01 | 0.090 | 0 | $1.54\times10^{-3}$ | 0 | 0 | $8.90\times10^{-4}$ | 0 | $8.90\times10^{-4}$ |
| 0.929 | $2.14\times10^{-3}$ | 0.04 | $-1.39\times10^{-2}$ | $1.13\times10^{-2}$ | $-1.21\times10^{-3}$ | $1.06\times10^{-5}$ | $3.21\times10^{-3}$ | $4.60\times10^{-5}$ | $8.48\times10^{-3}$ |
| 0.023 | 0 | 0.259 | -0.023 | 0 | 0 | 0 | 0 | 0 | 0 |
| **R pod** | | | | | | | | | |
| 0.548 | -0.01 | 0.090 | 0 | $1.54\times10^{-3}$ | 0 | 0 | $8.90\times10^{-4}$ | 0 | $8.90\times10^{-4}$ |
| 0.929 | $2.14\times10^{-3}$ | 0.04 | $1.39\times10^{-2}$ | $1.13\times10^{-2}$ | $-1.21\times10^{-3}$ | $1.06\times10^{-5}$ | $3.21\times10^{-3}$ | $4.60\times10^{-5}$ | $8.48\times10^{-3}$ |
| 0.023 | 0 | 0.259 | -0.023 | 0 | 0 | 0 | 0 | 0 | 0 |
| **LL pod** | | | | | | | | | |
| 0.571 | -0.01 | 0.091 | 0 | $1.54\times10^{-3}$ | 0 | 0 | $8.90\times10^{-4}$ | 0 | $8.90\times10^{-4}$ |
| 0.929 | $2.14\times10^{-3}$ | 0.04 | $-1.39\times10^{-2}$ | $1.13\times10^{-2}$ | $-1.21\times10^{-3}$ | $1.06\times10^{-5}$ | $3.21\times10^{-3}$ | $4.60\times10^{-5}$ | $8.48\times10^{-3}$ |
| 0.023 | 0 | 0.259 | -0.023 | 0 | 0 | 0 | 0 | 0 | 0 |
| **RR pod** | | | | | | | | | |
| 0.571 | -0.01 | 0.091 | 0 | $1.54\times10^{-3}$ | 0 | 0 | $8.90\times10^{-4}$ | 0 | $8.90\times10^{-4}$ |
| 0.929 | $2.14\times10^{-3}$ | 0.04 | $-1.39\times10^{-2}$ | $1.13\times10^{-2}$ | $-1.21\times10^{-3}$ | $1.06\times10^{-5}$ | $3.21\times10^{-3}$ | $4.60\times10^{-5}$ | $8.48\times10^{-3}$ |
| 0.023 | 0 | 0.259 | -0.023 | 0 | 0 | 0 | 0 | 0 | 0 |

**Table 3.8** Lumped mass data. Relative positions are given with respect to the pod/wing spar intersection node. Frame of reference: $x$ right wing, $y$ upstream, $z$ up.

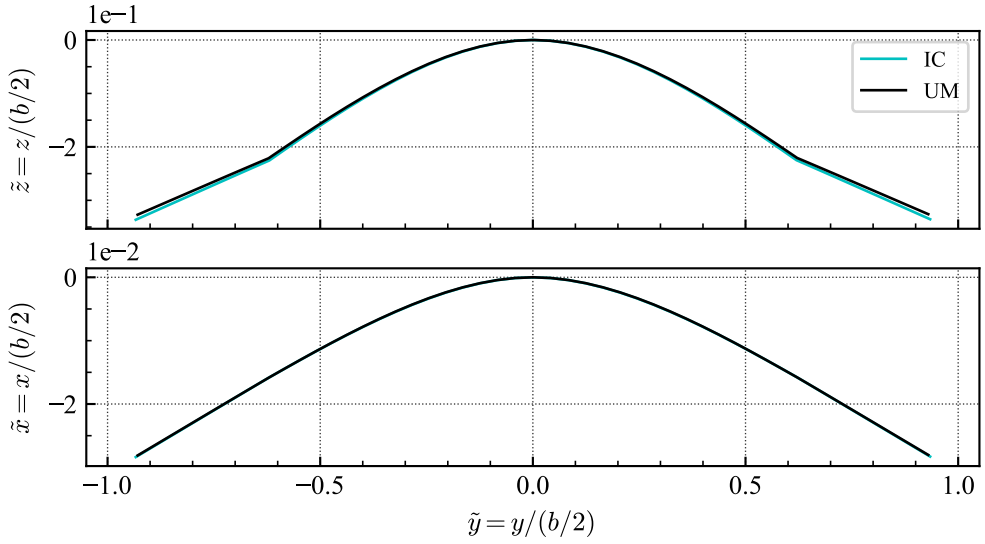



**Fig. 3.11** Spar deformation when subject to gravity.

| Part [-] | Span [m] | Chord [m] | E. Axis [% chord] | Dihedral [deg] | Airfoil [-] |
|---|---|---|---|---|---|
| inner/mid | 1 | 0.2 | 28.8 | 0 | EMX-07 |
| outer | 1 | 0.2 | 28.8 | 10 | EMX-07 |
| C tail | 0.385 | 0.11 | 32.35 | 0 | FLAT |
| LL/L/R/RR tail | 0.48 | 0.11 | 32.35 | 0 | FLAT |
| C fin | 0.15 | 0.78 | 122.56 | 0 | FLAT |
| L/C/R pod | 0.184 | 0.38 | 60.93 | 0 | FLAT |

**Table 3.9** X-HALE aerodynamic model description.

| Part | SHARPy | | UM/NAST | |
|---|---|---|---|---|
| [-] | M | N | M | N |
| inner/mid | 8 | 8 | 8 | 12 |
| outer | 8 | 16 | 8 | 12 |
| C tail | 3 | 4 | 8 | 6 |
| LL/L/R/RR tail | 3 | 4 | 8 | 6 |
| C fin | 4 | 2 | 8 | 2 |
| L/C/R pod | 4 | 2 | 8 | 2 |

**Table 3.10** Aerodynamic grid discretisation

30 N m). Follower-force effects are critical nonlinearities in aeroelasticity with large deformations, as the final orientation of the aerodynamic loads depends on the beam displacements and rotations at the equilibrium point. The wing spar nodal positions for both load cases and solvers are presented in Figure 3.12. A relative deflection of about 35% of the semispan is reached with the given inputs. As it can be seen, excellent comparison is obtained between both structural solvers even under such extreme loading conditions.

Lastly, the structural modes and natural frequencies of the aeroelastically trimmed structure have been extracted with both codes. The results up to the fifth mode are presented in table 3.11. The small differences are due to the different methods used here and their convergence criteria not being exactly equivalent.

### 3.4.3 Aerodynamic cases

The total aerodynamic forces in wind-oriented frame of reference have been calculated for the underformed main wing. In order to assess the VLM results isolated from the aeroelastic effects and the sensitivity of the inclusion of a non-symmetrical airfoil, the resultant forces for a symmetric and the cambered airfoil EMX-07 (Fig. 3.13) in

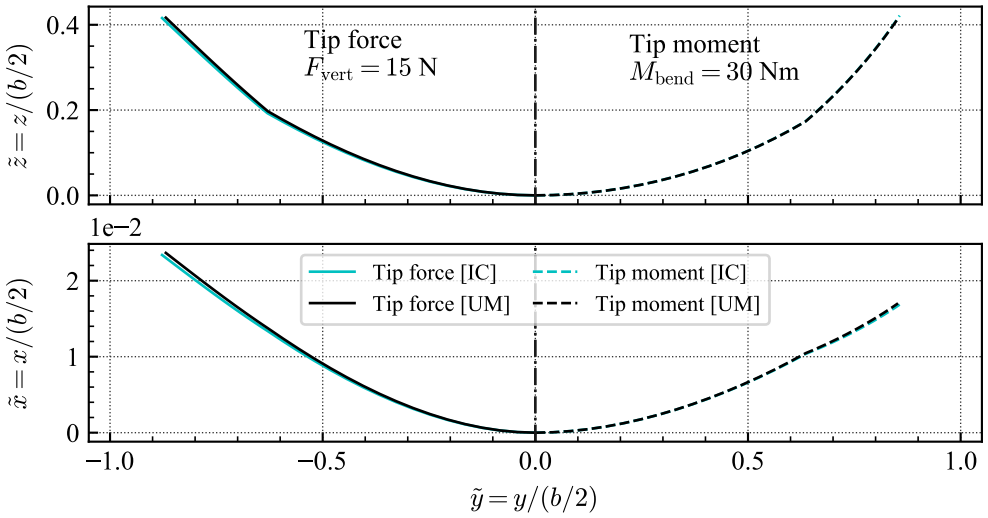



**Fig. 3.12** Spar nodal positions under a follower force ($\tilde{y} < 0$) and moment ($\tilde{y} > 0$).

| Modes | Mode type | Freq. - IC | Freq. - UM |
|---|---|---|---|
| 1 & 2 | Out of plane I | 0.57 | 0.59 |
| 3 & 4 | Torsion I | 2.50 | 2.57 |
| 5 & 6 | Out of plane II | 3.62 | 3.70 |
| 7 & 8 | In plane I | 4.58 | 4.45 |
| 9 & 10 | Torsion II | 6.59 | 6.57 |

**Table 3.11** Structural natural frequencies (in Hz) of the aeroelastically trimmed X-HALE structural model.

the main wing are calculated. As an external reference, the same geometry has been simulated in TORNADO [114], which is an open-source steady vortex-lattice solver with a horseshoe wake. The distribution of aerodynamic forces plays a very important role in the final aeroelastically trimmed shape. Studies in this work have shown a strong sensitivity of the distributed aerodynamic forces and beam deflections in aeroelastic equilibrium with changes in the camber line. Therefore, special care must be taken when calculating and discretising the camber line from the airfoil coordinates when performing nonlinear aeroelastic analysis using lifting surface models. In the studies conducted by the authors, it has been determined that for static simulations with X-HALE, 8 or more chordwise panels distributed uniformly replicate the camber line satisfactorily. This strong dependency of the results with the camber line shape is also evident when considering the importance of the correctness of the camber line determination. Determining the camber line from a set of airfoil coordinates is not a trivial exercise. Decomposing the

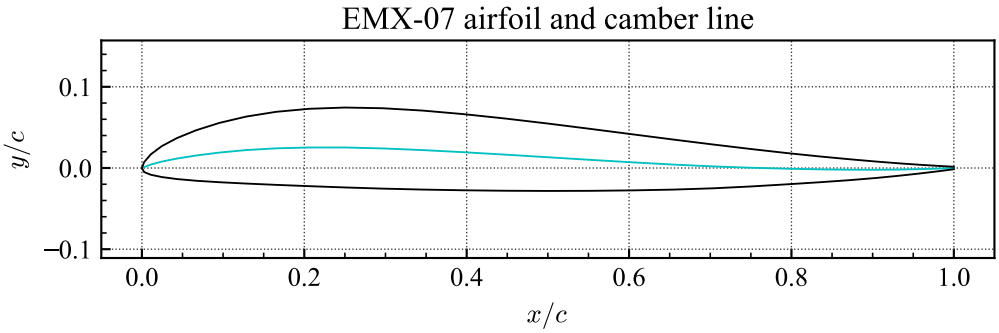



**Fig. 3.13** EMX-07 airfoil used in main wing.

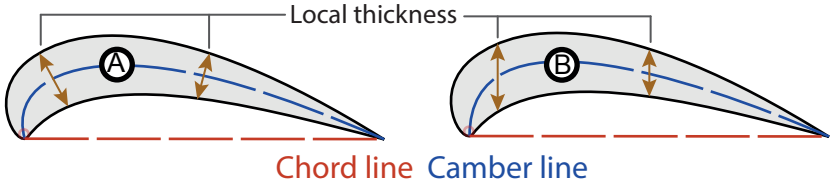



**Fig. 3.14** Difference between the American (A) and the British (B) camber definition.

airfoil coordinates as:

$$z_{\text{upper}}(x) = z_{\text{camber}}(x) + \tfrac{1}{2}t(x) \tag{3.2}$$

$$z_{\text{lower}}(x) = z_{\text{camber}}(x) - \tfrac{1}{2}t(x), \tag{3.3}$$

the thickness distribution $t(x)$ needs to be computed in order to obtain the camber distribution $z_{\text{camber}}$. There are two conventions concerning the definition of the thickness (and by extension, of the camber line): the "*American convention*" and the "*British convention*". The former is featured in [77, 139]. It defines the airfoil thickness as the perpendicular *to the camber line* distance between the airfoil sides. In contrast, the British convention [11] measures the thickness as the perpendicular *to the chord line* distance. Figure 3.14 illustrates the difference between the American convention (A) and the British convention (B). This work relies on the American definition. Although it is more difficult to compute with this interpretation, the results are more accurate than when the British definition is compared against a solver with thickness support. The difference between these two methods tends to none with small curvature airfoils. However, during this study, it has been found that for a reflex airfoil such as the EMX-07, the method used for obtaining the camber line can yield discrepancies in the lift and drag

| Code | Airfoil | Induced Drag [N] | Lift [N] |
|---|---|---|---|
| IC-SHARPy | NACA 0012 | 0.28 | 58.60 |
| | EMX-07 | 0.33 | 63.61 |
| UM-NAST | NACA 0012 | 0.29 | 61.69 |
| | EMX-07 | 0.32 | 64.00 |
| TORNADO | NACA 0012 | 0.28 | 58.80 |
| | EMX-07 | 0.31 | 61.46 |

**Table 3.12** Aerodynamic forces on the rigid aircraft, $\alpha = 4°$.

| | |
|---|---|
| Free stream velocity, $U_\infty$ | 14 [m/s] |
| Density, $\rho_\infty$ | 1.225 [kg/m$^3$] |
| Gravity acceleration, $g$ | 9.807 [m/s$^2$] |

**Table 3.13** Flight conditions used for the X-HALE validation cases.

predictions of almost 10%, which once they are introduced in a coupled environment, can become even larger.

Table 3.12 shows the total aerodynamic forces in wind axes for a wing angle of attack $\alpha = 4°$, and Table 3.13 shows the flight conditions used for the results presented here. Induced drag prediction in VLM codes is more difficult to achieve than lift due to their quadratic dependency and the smaller order of magnitude of the former. Due to the potential flow assumption, viscous drag needs additional information such as Reynolds number-dependant $C_D$ *vs* $C_L$ databases from experimental or numerical methods or other empirical laws. In sight of the sensitivity of the results with the camber line input, the results shown in Table 3.12 are considered to be satisfactory.

### 3.4.4 Aeroelastic cases

The aeroelastic cases presented in this section are of two kinds. The first is a comparison of the structural deformations at a given flight condition (angle of attack, tail deflection, flight speed, freestream density) for both codes while the second is a comparison of the trim conditions calculated with each code.

The aeroelastic deformation comparison given a flight condition has been run for three different body angles of attack: 0, 3 and 5 degrees. It is important to note that the body angle of attack is defined as the angle between the free stream and the undeformed central fuselage, as the main wing presents a structural twist (and thus, an offset angle of attack) of 5 degrees. In addition, the thrust has been set to zero and the tails are aligned with the fuselage ($\delta = 0°$). The aircraft is assumed to be clamped on the central

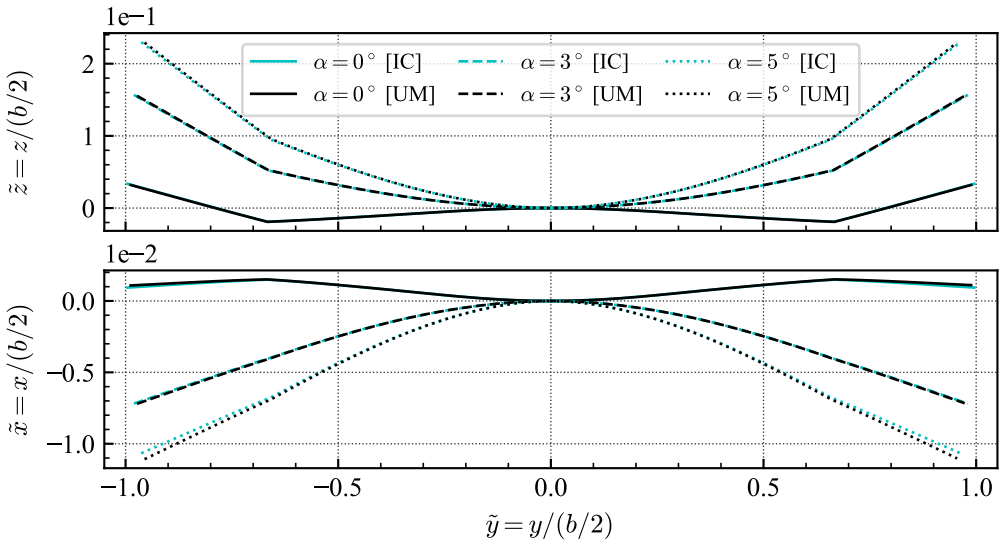



**Fig. 3.15** Spar deformation for different angles of attack

fuselage-main spar intersection. Results are shown in Figure 3.15. From it, one can see that even in deflections larger than 20% of the semi wing span, the difference in wing deformation between both solvers is almost negligible, showing again their suitability for this kind of problem.

Trimming a very flexible aircraft is more complex than a rigid one. In addition, the unconventional configuration of the X-HALE contributes to this difficulty. Given the relatively low torsional stiffness of the wing spar, the effect of the multiple tails is not straightforward: an increase in control surface deflection ($\delta > 0$) has two effects. First, a nose-down pitching moment and a momentary contribution towards increasing the overall lift; and second, a loss of lift in the main wing, specially in the outer parts of the wing due to a torsion load that reduces the effective twist. The second contribution is what is classically identified as "control reversal".

X-HALE has some characteristics that, if considered in the trim routine, can simplify the analysis. First, the four outer tails are actuated using a single input channel. That means that a single "tail deflection" ($\delta$) parameter can be used. Second, the five motors are controlled using a *symmetric* and *antisymmetric* thrust in order to provide lateral control. Given that the only non-symmetric contribution comes from the mass distribution, and the centre of gravity is deviated less than 4 mm from the longitudinal axis, trim is sought using the longitudinal equations only. The residual forces and moments in the lateral equations after trim are all the same order of magnitude (or lower) than the trim

| **Code** | **Tail orientation** | $\boldsymbol{\alpha}$ [°] | $\boldsymbol{\delta}$ [°] | $\boldsymbol{T}$ [N] |
|---|---|---|---|---|
| IC-SHARPy | Vertical | 2.64 | 1.19 | 0.223 |
| | Horizontal | 2.21 | 0.52 | 0.213 |
| UM-NAST | Vertical | 2.59 | 1.15 | 0.179 |
| | Horizontal | 2.38 | 0.66 | 0.179 |

**Table 3.14** Trim results for cruise flight

tolerances, and consequently, the *antisymmetric* component of the thrust vector is set to zero in forward flight conditions.

In order to enforce straight and steady unperturbed flight in relatively long dynamic simulations, the total residuals from the trim routine are still applied (with opposite sign) at the wing spar centre as a follower force. It is important to note however that applying those small constant balancing loads will not affect the vehicle dynamics, which is governed by loads variations from the equilibrium condition.

The computed trim conditions for both the vertical and horizontal central tail orientations are presented in Table 3.14. The agreement for $\alpha$ and $\delta$ between codes is good, while the thrust reflects the difference in drag prediction. Being the drag counteracted by the span-distributed propellers, the effect of the drag discrepancy on the flight dynamics will not affect the results between codes.

## 3.5 X-HALE gust response

The open-loop response of the X-HALE has been simulated with SHARPy and UM/NAST and compared on deterministic "1-cos" gust profiles. In order to evaluate the response of the aircraft to the gust, trajectory, attitude angles, wing root loads and wing tip deflection are recorded. All results shown here are given in inertial, $G$, or local material, $B$ frame of reference. It is important to note then, that the plane is flying towards a negative $x$ and consequently, an increase in forward speed translates in $\Delta V_x < 0$.

An array of different gusts has been defined and the response of the X-HALE with vertical and horizontal central tail has been simulated. The analysis has been simplified by setting the gust reference velocity ($\bar{U}_g$) to 15% of the free stream speed. It has been chosen so that the surfaces only operate in the linear regime, as a coarse approximation for instantaneous angle of attack is $\alpha = \arctan 0.15 = 8.53°$. The "1-cos" velocity profile

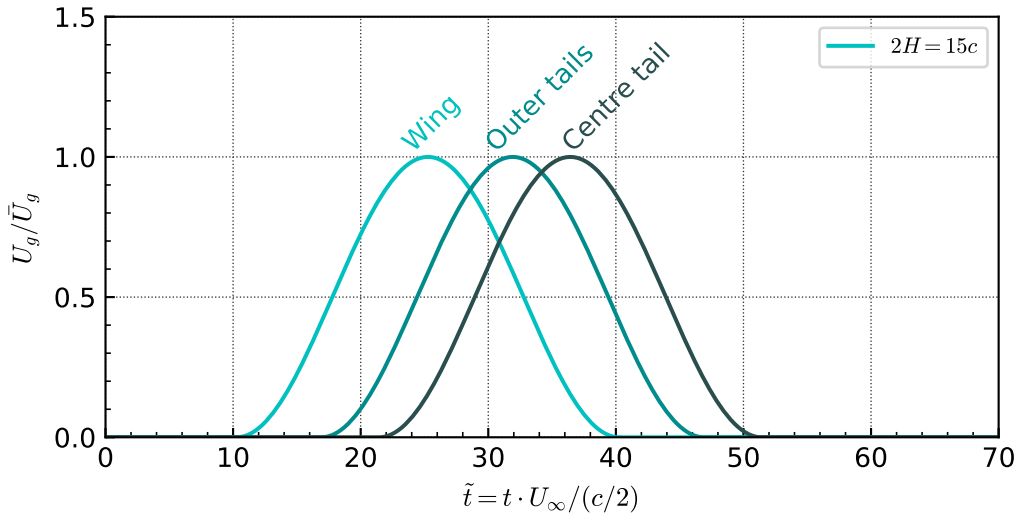



**Fig. 3.16** Gust velocity time histories for different aerodynamic surfaces (as seen from the elastic axis in steady forward flight).

is given as

$$U_g = \frac{\bar{U}_g}{2}\left(1 - \cos\left(\frac{x - x_{\text{offset}}}{H}\pi\right)\right) \; ; \qquad \text{if } -2H < x - x_{\text{offset}} < 0 \tag{3.4}$$

where $H$ is the gradient of the gust, that is, the distance required to reach the reference velocity and $x_{\text{offset}}$ is the initial position of the gust front with respect to the aircraft main spar.

Given the number of lifting surfaces and their different locations, it is important to know when the gust is reaching each of the different tails and fins. Figure 3.16 presents the time history of the gust contribution to the incoming velocity for several lifting surfaces with relation to the reference axis of each surface, such as the main wing, the outer tails and the centre tail. In this plot, and in all the remaining ones, the time variable is nondimensionalised as

$$\tilde{t} = t \cdot U_\infty/(c/2). \tag{3.5}$$

This choice of a convective definition of the nondimensional time is a reflection of the focus of this paper in a short period after the gust, where structural response and flight dynamics are tightly coupled. The long-term effects of the gust, where the flight dynamics

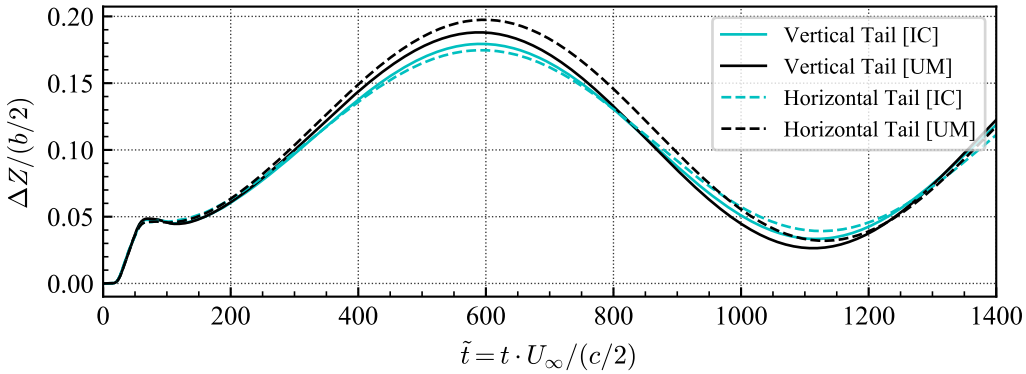



**Fig. 3.17** $2H = 15$ chords, vertical gust response: wing spar center trajectory.

and the flight path in general can present significant deviations from the undisturbed case, and are highly dependent on the control surface inputs, are left for future work.

It is also important to note that all results presented here include a $x_{\text{offset}} = -1\,\text{m}$ as defined in eq. (3.4). The simulations have been run for 10 physical seconds, which equates to $\tilde{t}_{\text{max}} = 1400$.

The wing spar root vertical trajectory is shown in Figure 3.17 for the vertical and horizontal tail configurations. It is easy to observe an oscillatory motion in the longitudinal plane. An interesting observation is that both codes agree that the period and amplitude of this mode is weakly dependent on the orientation of the central tail. Figure 3.18 shows the trajectory velocity in body frame of reference. The agreement between codes is good, specially in the transient phase of the response, which is the most relevant for loads. The $V_x$ component of the velocity (upper plot) presents a feature that will be present in the majority of results here shown: the results corresponding to the SHARPy code present lower damping. In this case, the oscillations in $V_x$ correspond to the in-plane bending motion of the wing that transmit to the wing root. This would not be found on the corresponding velocity component at the centre of gravity. However, given that for very flexible aircraft the centre of gravity position is dependant on the structural deformations, it has been decided to use a fixed point in body frame of reference. Figure 3.18 shows the rigid body velocity in body-attached ($B$) frame of reference and it is better understood together with the pitch angle increment history (Fig. 3.19). Figure 3.16 shows that the outer tails enter the gust around $\tilde{t} \approx 20$ and reach the peak velocity at $\tilde{t} = 35$. This can be seen in Figure 3.19, where the pitch reaches a maximum due to the wing-only excess of lift and then the increasingly stronger tail influence causes it to pitch down heavily. The influence of the main wing in the pitch up part of the

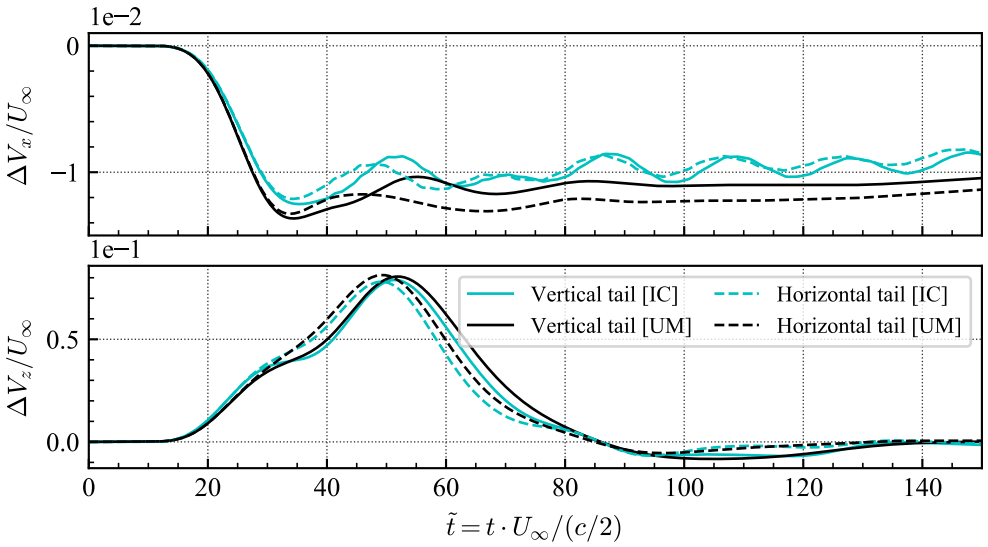



**Fig. 3.18** $2H = 15$ chords, vertical gust response: body velocity in body-attached frame of reference. Long simulation (10 s in physical time).

response is weaker than could be expected from conventional aircraft. This is due to the span-loading character of the X-HALE, where the mass is distributed along the span, rather than the fuselage, like typical aircraft do. Because the main contribution to the lift and mass of the aircraft comes from the main wing, the aerodynamic centre and the centre of mass are located close to each other. Thus, a lift increment in the wing causes a vertical velocity increase, but not a strong pitch up tendency. The fact that the centre of lift and gravity are so closely located also affects the damping of the longitudinal rigid-body motion.

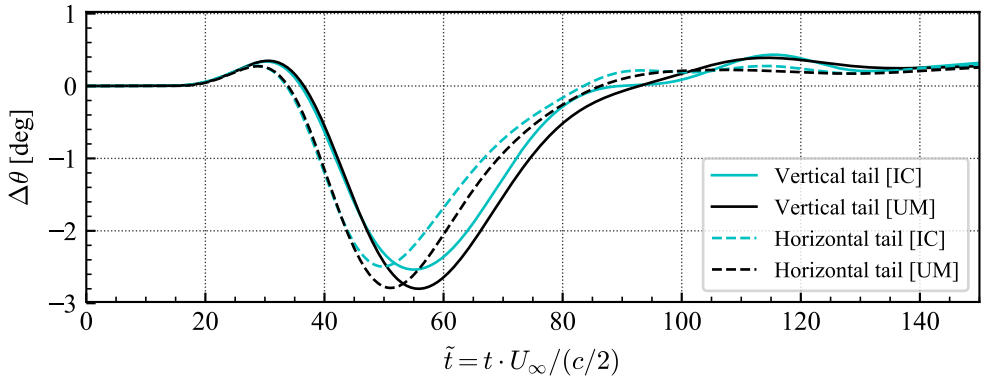



**Fig. 3.19** $2H = 15$ chords, vertical gust response: wing root pitch angle history.

The tail orientation also has an effect on the pitch down part of the response ($30 \lesssim \tilde{t} \lesssim 50$). The pitch angle peak value is similar for both tail orientations. In contrast, the peak is reached earlier in time for the horizontal tail. This is a predictable result, as the increased horizontal wet area in the aft part of the aircraft will make the aircraft pitching characteristics more responsive to perturbations in $\alpha$.

Other interesting feature of the longitudinal response is the lowly damped phugoid motion. A rough approximation for the angular frequency and damping of the phugoid mode can be obtained from classic flight dynamics theory with a few other simplifications such as $V_z = \dot{V}_z = 0$ and $\alpha = 0$ [50]. The resulting expressions for the phugoid angular frequency, $\omega_n$, and damping, $\zeta$, are:

$$\omega_n = \sqrt{2}\frac{g}{U_0}, \quad \zeta = \frac{1}{\sqrt{2}}\frac{C_{D0}}{C_{L0}} \tag{3.6}$$

where $U_0 = U_\infty = 14\,\mathrm{m\,s^{-1}}$ and $C_{L0}/C_{D0}$ is the efficiency in cruise conditions. Substituting in the angular frequency equation, we obtain that the period for the phugoid given by this approximation is $T_{\text{phugoid}} = (\omega_n/(2\pi))^{-1} = 6.34$ s. The longitudinal oscillatory mode in Figure 3.17 has a period of 7.5 s. This is a reasonably good agreement in sight of the coarseness of the approximation leading to equation (3.6), where no centre of mass position, $C_{mq}$ or $C_{m\dot{\alpha}}$ influence is considered, as the pitching-moment equation was not considered.

The approximation for the damping in equation (3.6) effectively states that the more efficient an aircraft is, the more undamped the phugoid mode. Given the high aspect ratio of the X-HALE, its efficiency is also high, and thus $C_{D0} \ll C_{L0}$. The resulting damping coefficient will be small in first approximation. Other simulations conducted by the authors show an influence of the central tail orientation in the damping of the motion. This behaviour cannot be seen with the analytical approximations used here, as the influence of the tail would be captured through $C_{mq}$ and $C_{m\dot{\alpha}}$. The cases here shown have not been run long enough to extract a definite trend from the trajectory about the damping of the phugoid, but in sight of the first cycle, the horizontal tail case shows a first tendency to damp the motion quicker than the vertical tail one.

Lastly, the time history of the wing tip deformation for the early stages of the response is shown in Figure 3.20. Here, the SHARPy deformation history of the wing in the in-plane motion shows an oscillatory motion that could also be seen in the $V_x$ component of the body velocity in Figure 3.18. We note that in-plane structural vibrations are challenging to determine accurately, specially when running potential-flow solvers, where

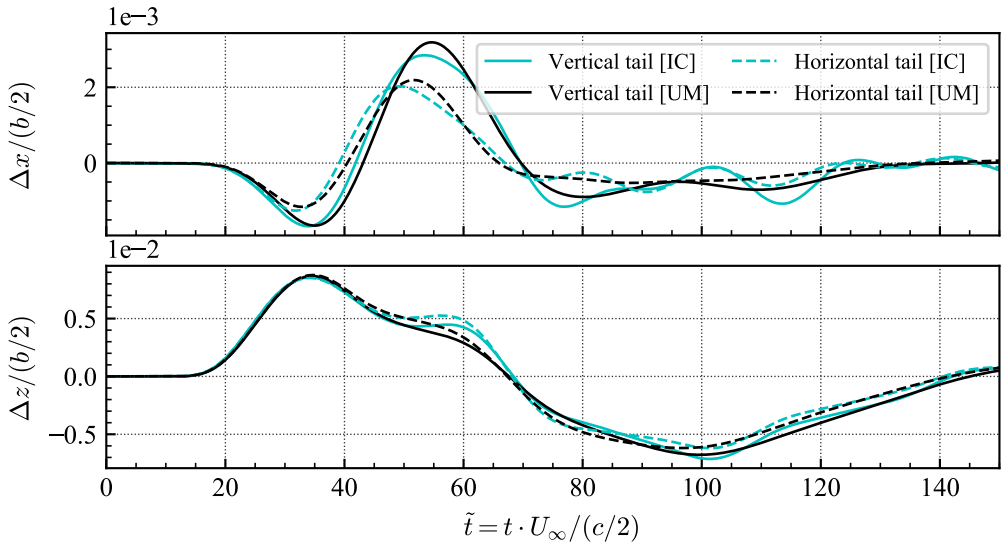



**Fig. 3.20** $2H = 15$ chords, vertical gust response: wing tip displacement in body-attached axes.

the main contribution to in-plane force is the induced drag. The oscillations in this direction have an effect in the effective incoming speed of the flow, and this effect is much weaker than the out-of-plane equivalent, where aerodynamic damping plays a major role in the aeroelastic response. On the other hand, vertical displacements, which are an order of magnitude larger and directly affected by the input vertical gust, show an excellent agreement.

This chapter is complemented by Ref. [40], which deals with the lateral gust response of the X-HALE configuration, including discussions related to the discretisation effect on the accuracy of the solution for this kind of simulation.

## 3.6 The T-tail HALE model

The T-tail HALE configuration has been designed in order to provide a simple, easily reproducible, conventional HALE example. It is based on the straight wing used in [174] and its extension to a T-tail configuration by [69]. It features a high aspect ratio wing with very low stiffness in torsion and low out of plane bending stiffness. The tail is arranged in a T-configuration, with elevator and rudder as control surfaces. All the aircraft structure is modelled as flexible, although the fuselage and tail are significantly stiffer (10×) than the wing in order to simplify the analysis of the flight dynamics. All the surfaces have symmetric airfoils, thus they are modelled as flat plates in the UVLM.

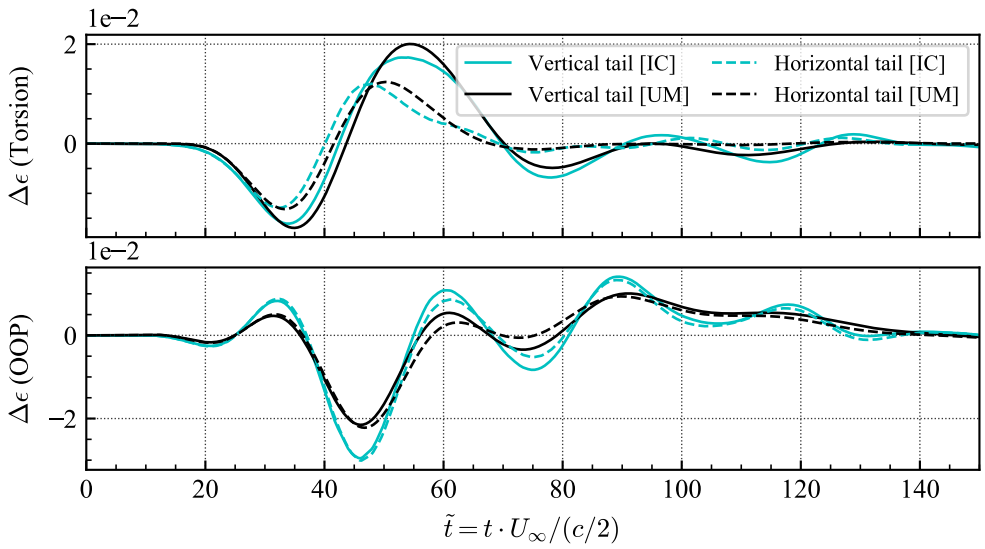



**Fig. 3.21** $2H = 15$ chords, vertical gust response: right wing root element strains.

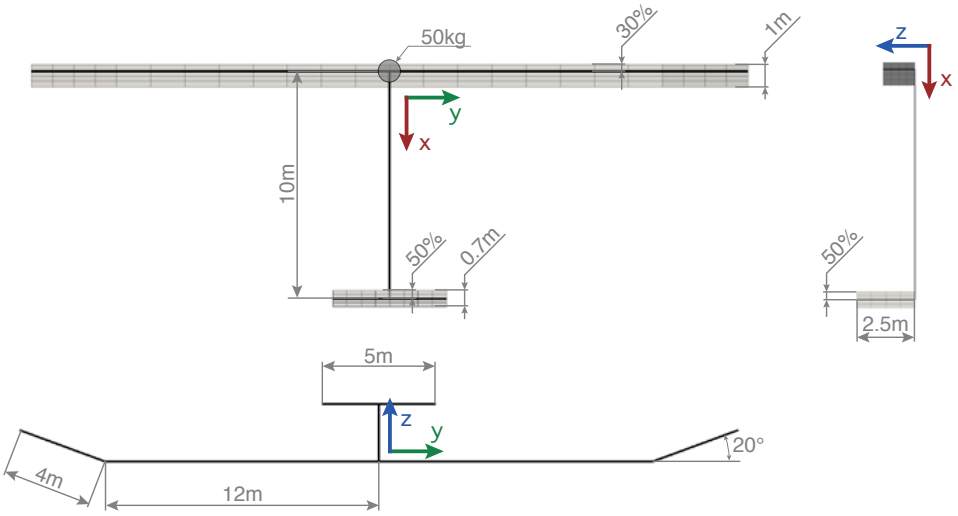



**Fig. 3.22** Dimensions of the HALE T-tail configuration.

| **Component** | **EA** | **GA$_y$** | **GA$_z$** | **GJ** | **EI$_y$** | **EI$_z$** | $\bar{\mathbf{m}}$ | $\bar{\mathbf{J}}$ |
|---|---|---|---|---|---|---|---|---|
| Wing | 1.5e6 $\cdot\sigma$ | 1.5e6 $\cdot\sigma$ | 1.5e6 $\cdot\sigma$ | 1.5e4 $\cdot\sigma$ | 3.0e4 $\cdot\sigma$ | 6.0e6 $\cdot\sigma$ | 0.75 | 0.1 |
| Fuselage | 1.5e7 $\cdot\sigma$ | 1.5e7 $\cdot\sigma$ | 1.5e7 $\cdot\sigma$ | 1.5e5 $\cdot\sigma$ | 3.0e5 $\cdot\sigma$ | 6.0e7 $\cdot\sigma$ | 0.1 | 0.01 |
| Tail | 1.5e7 $\cdot\sigma$ | 1.5e7 $\cdot\sigma$ | 1.5e7 $\cdot\sigma$ | 1.5e5 $\cdot\sigma$ | 3.0e5 $\cdot\sigma$ | 6.0e7 $\cdot\sigma$ | 0.1 | 0.01 |

**Table 3.15** Stiffness and mass data for the HALE model.

| Variable | Flexible | Very Stiff |
|---|---|---|
| Angle of attack, $\alpha$ [°] | 4.072 | 3.81 |
| Control surface deflection, $\delta$ [°] | -1.27 | -0.76 |
| Rudder deflection, $\delta_r$ [°] | 0.00 | 0.00 |
| Thrust, $T$ [N] | 3.87 | 3.76 |
| Total mass [kg] | 78.25 | |
| Wing surface [m$^2$] | 32.00 | |
| Wing loading [kg m$^{-2}$] | 2.44 | |

**Table 3.16** Trim values for the baseline (flexible) and stiff T-tail HALE model at cruise speed ($10\,\mathrm{m\,s^{-1}}$).

Figure 3.22 provides an overview of the geometry of the HALE model. It can be seen that the only lumped mass is located at the spar-fuselage intersection, weights 50 kg (75% of the total mass) and it is modelled as a point mass, with no moment of inertia. This aircraft features two control surfaces: an elevator and a rudder. The rudder has a chord of 25% of the tail chord, while the elevator is a moving surface pivoting around the elastic axis. Propeller force is modelled as a follower force applied in the spar-fuselage intersection. Lastly, table 3.15 shows the mass and stiffness data. The baseline case is run with the stiffness multiplier, $\sigma = 1.5$. $1g$ trim values for the baseline and the stiff —$\sigma = 50$— configurations at cruise speed are shown in Table 3.16.

The element mass $\boldsymbol{\mathcal{M}}$ and stiffness $\boldsymbol{\mathcal{K}}$ matrices are assembled as follows:

$$\boldsymbol{\mathcal{M}} = \begin{bmatrix} \bar{m}\mathbb{I}_{3\times 3} & -\bar{m}\tilde{\boldsymbol{\xi}}_{\mathrm{cg}} \\ \bar{m}\tilde{\boldsymbol{\xi}}_{\mathrm{cg}} & \mathrm{diag}\left(\bar{J}, \bar{J}/2, \bar{J}/2\right) \end{bmatrix}; \qquad \boldsymbol{\mathcal{K}} = \mathrm{diag}(EA, GA_y, GA_z, GJ, EI_y, EI_z), \tag{3.7}$$

where $\boldsymbol{\xi}_{\mathrm{cg}}$ is the cross-sectional coordinates of a point with respect to the elastic axis. These are given in material frame of reference (noted as $B$ in figure 2.1). The stiffness properties are as follows: $EA$ the axial stiffness, $GJ$ the torsional stiffness, $EI_y$ the out-of-plane bending stiffness, and $EI_z$ the in-plane stiffness. The shear stiffness, $GA_y$ and $GA_z$, have been chosen high enough to dismiss the effect of these degrees of freedom (in this case, $1 \times 10^5\,\mathrm{N\,m^2}$), but low enough so that the numerical stiffness of the structural problem is kept manageable. The cross-sectional moments of inertia in the mass matrix $\mathrm{diag}\left(I_{xx}, I_{yy}, I_{zz}\right) = \left(\bar{J}, \bar{J}/2, \bar{J}/2\right)$ have been chosen such that they comply with $I_{xx} = I_{yy} + I_{zz}$. This configuration presents large wing deformation even in steady horizontal flight ($\sim 24\%$). Due to the effective dihedral, in-plane and torsion loads are coupled.

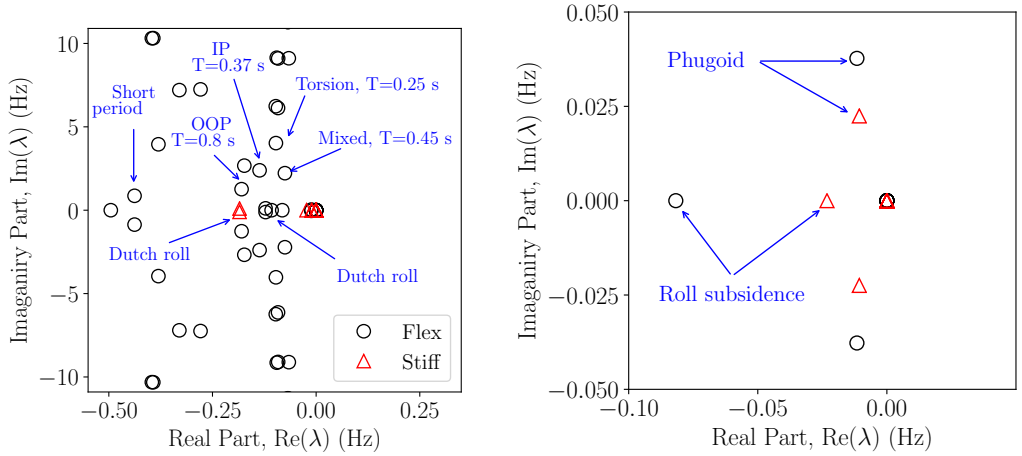



**Fig. 3.23** Dynamic stability diagram for rigid and flexible aircraft. Right: Eigenvalues in the physical region of interest, left: zoomed-in plot for flight-dynamic modes.

Figure 3.23 shows the stability plot near the origin for both stiff ($\sigma = 50$) and flexible ($\sigma = 1.5$) aircraft in trimmed flight. It displays the smaller eigenvalues of the linearised vehicle dynamics around a nonlinear static equilibrium. It can be seen that both aircraft are asymptotically stable (real parts of all modes are negative) and therefore define a suitable starting point for the following nonlinear dynamic analysis. The zoomed-in (right-hand side) part of figure 3.23 focuses on the flight-dynamic modes of both variants. For example, the phugoid can be seen at a slightly different natural frequency than expected for larger scale aircraft. Other modes observed are the short period and the lateral-directional modes such as the Dutch roll and roll divergence ones. The zoomed-out version of the eigenvalue plot (left-hand side) also includes the lower-frequency aeroelastic modes on the flexible aircraft. Very different levels of damping can be observed in the elastic modes of this vehicle, as out-of-plane wing displacements have much higher aerodynamic damping than in-plane bending and torsional dynamics [69].

# Chapter 4

# Low altitude atmospheric turbulence impact on aircraft aeroelasticity

In order to investigate the sensitivity of the different methods used to obtain a continuous turbulent field for flexible aircraft atmospheric disturbance loads, we examine the response of a representative very flexible aircraft model to three atmospheric turbulence input methods namely: the one-dimensional von Kármán, a two-dimensional stochastic formulation: the Kaimal model [88] enhanced with a spatial coherence model, as well as full-three dimensional wind fields produced by means of large-eddy simulations [141, 15, 42]. This chapter starts with a description of the setups used to simulate the wind fields, together with comparing their statistics and respective velocity spectra. Subsequently, SHARPy is used to make time-domain predictions of the flight dynamics and loads for a representative HALE configuration with very flexible structure and slender wings, which will be described in this chapter. Results from the simulations including a sensitivity analysis for the three turbulence inputs and a cross-comparison between a flexible and a stiffer aircraft are presented in Section 4.4.

## 4.1 Atmospheric turbulence modelling

Within the atmospheric boundary layer (ABL), turbulence evolves in response to the local topography as well as the heating and cooling of the Earth's surface. The latter, categorises the ABL in three distinct states, namely, a neutral boundary layer (NBL), a convective boundary layer (CBL) and a stable boundary layer (SBL). In the neutral boundary layer, thermal effects remain unimportant and air parcels move adiabatically

solely due to Coriolis forcing and surface friction. Thermal effects on the other hand are significant in the convective and stable boundary layers in which temperature stratification leads to a capping inversion (inversion layer). In the CBL, capping inversion acts as a lid and turbulence convection is dominated by large plumes (thermals) that transport and mix air away from the surface. The CBL state is often found following the sunrise. On the other hand, the SBL is characterised by a strong temperature stratification which typically leads to a shallower boundary layer thickness, strong wind shear and much smaller and weaker eddies than the neutral and convective cases [208]. Nonetheless, all three states (except for the convective case) are very rarely in a state of equilibrium and transition from one to another takes place multiple times during the day. However, provided the disparity between the ABL state evolution ($\sim$ hours) and the time spent by a VFA in the lower ABL ($\sim$ minutes), the assumption of ABL equilibrium can be considered to be valid. Assuming that the ABL under consideration to be in a state of equilibrium allows us to also assume the atmospheric turbulence stationary character and therefore to describe it statistically (e.g. through its mean speed and turbulence intensity). In simulating atmospheric turbulence, two approaches can be generally considered. The first approach uses stochastic formulations to generate wind fields based on first and second-order statistics (spectral/co-spectra). Stochastic generation of wind fields has been an active area of research since the 1950s [76, 194, 105, 106]. For applications in wind engineering [82, 60], where an relatively accurate low altitude ABL model is important, the most common stochastic model is that of Kaimal [88, 194]. Aircraft loads arising from atmospheric disturbances were computed with discrete gust profiles such as step and 1-cos up until the late 1960s, when continuous turbulence spectra were defined and the theoretical basis for the analysis were mature enough [72]. Most of continuous turbulence analysis for aircraft loads have been conducted with Dryden or von Kármán spectra, both of them being span-constant approximations. The second technique on the other hand, utilises detailed numerical simulations to resolve the fluid flow equations (or a parametrisation of them) for a given temporal and spatial scale. Such a popular method is large-eddy simulation (LES). The idea underpinning LES lies in resolving the large energetic structures of the flow while modelling the smaller scales. LES produces three-dimensional wind field for a given computational domain and its quality depends on the resolution and modelling parameters. To this end, LES does not resolve all relevant turbulent scales, but it provides (when properly designed) excellent agreement with theoretical model such as the Monin-Obhukov similarity theory (MOST) [121] and experimental data [68]. Further details on both approaches are discussed next.

### 4.1.1 Stochastic formulations for turbulence generation

In analysing the dynamic response of an aircraft, three-dimensional isotropic turbulence models are often used to generate the turbulence input. They produce stochastic time signals for the three velocity components $(u, v, w)$ by using either the Dryden or the von Kármán spectrum. The three-component signal is produced by passing band-limited white noise through filters obtained from spectral factorisation of the semi-empirical models [104, 17]. Uniform signals in transversal planes are subsequently generated in the area of interest creating a velocity field with perfect streamwise coherence, $\mathrm{coh}_{uu} = 1$. In this study we use the von Kármán model, which is widely accepted as more accurate than the Dryden one. For a nominal reference velocity $U_{ref}$, it is defined by the power spectral density,

$$S_u(f) = \frac{4\sigma_u^2 L_u}{U_{\mathrm{ref}}} \cdot \frac{1}{\left[1 + \left(1.339 L_u \frac{2\pi f}{U_{\mathrm{ref}}}\right)^2\right]^{5/6}}, \tag{4.1}$$

for the longitudinal velocity and

$$S_{v,w}(f) = \frac{4\sigma_{v,w}^2 L_{v,w}}{U_{\mathrm{ref}}} \cdot \frac{1 + \frac{8}{3}\left(2.678 L_{v,w} \frac{2\pi f}{U_{\mathrm{ref}}}\right)^2}{\left[1 + \left(2.678 L_{v,w} \frac{2\pi f}{U_{\mathrm{ref}}}\right)^2\right]^{11/6}}, \tag{4.2}$$

for the lateral velocities. As our interest is in the low-altitudes, independent turbulence length scales $L_u, L_v, L_w$ are assumed in three directions. They are chosen as in [115] so that $2L_w = h$ and $L_u = 2L_v = h/(0.177 + 0.000823h)^{1.2}$, where $h$ is the altitude (in feet). The corresponding turbulence intensities are $\sigma_u, \sigma_v, \sigma_w$. For the simulations, a time-signal of 250 s is generated and subsequently convected into a full 3-D wind field at speed $U_{ref}$ by means of "frozen turbulence", $\boldsymbol{u}(\boldsymbol{x}, t) = \boldsymbol{u}(\boldsymbol{0}, x - tU)$.

A more accurate representation of low-altitude atmospheric turbulence can also be sought via the Kaimal model. The Kaimal model can be used to make predictions of the wind field during the neutral boundary layer state using empirically calibrated constants. In the present study, Kaimal is used in conjunction with a spatial coherence model for the streamwise component of the velocity field, $u(y, z, t)$, known as the International Electrotechnical Commission (IEC) exponential coherence model [81],

$$\mathrm{coh}_{uu}(f, r, U_{\mathrm{ref}}) = \exp\left[-a\sqrt{\left(\frac{rf}{U_{\mathrm{ref}}}\right)^2 + \left(\frac{rb}{L_u}\right)^2}\right], \tag{4.3}$$

where $a = 10$ and $b = 0.12$ are the coherence decrement and offset parameters, $L_u$ the streamwise turbulence length scale, $U_{\text{ref}}$ is the mean wind speed at the reference height, $f$ denotes frequency and $r$ is a radial distance in a plane orthogonal to the streamwise direction. To generate the 2D wind field we make use of the open-source code TurbSim [87]. As TurbSim outputs planar fields of the velocity field as a function of time we use again Taylor's frozen turbulence hypothesis [184] to generate a full three-dimensional velocity field.

### 4.1.2 High-fidelity turbulent flows via large-eddy simulation

Large-eddy simulations (LES) of atmospheric turbulence are conducted using our in-house solver WInc3D [43]. It is an explicit LES framework which solves the unsteady, incompressible, filtered Navier-Stokes equations using their skew-symmetric form

$$\frac{\partial \widetilde{u}_i}{\partial t} + \frac{1}{2}\left(\widetilde{u}_j \frac{\partial \widetilde{u}_i}{\partial x_j} + \frac{\partial \widetilde{u}_i \widetilde{u}_j}{\partial x_j}\right) = -\frac{1}{\rho}\frac{\partial \widetilde{p^*}}{\partial x_i} - \frac{1}{\rho}\frac{\partial p_\infty}{\partial x_i} - \frac{\partial \tau_{ij}}{\partial x_j}, \quad i,j = \{1,2,3\}, \tag{4.4}$$

$$\frac{\partial \widetilde{u}_i}{\partial x_i} = 0, \tag{4.5}$$

where $\widetilde{p^*} = \widetilde{p} + 1/3\widetilde{u}_i\widetilde{u}_i$, and $\widetilde{u}_i$ are the filtered components of the modified pressure and velocity fields, respectively, $\rho$ is the fluid density, and indices $(1,2,3)$ correspond to a $(x,y,z)$ Cartesian coordinate system. The sub-filter stresses $-\partial_j \tau_{ij}$, that is, the residual terms from the filtering process, are calculated using the standard Smagorinsky model [172]

$$\tau_{ij} = -2(C_S\Delta)^2 |\widetilde{S}| \widetilde{S}_{ij}, \quad \widetilde{S}_{ij} = \frac{1}{2}\left(\frac{\partial \widetilde{u}_i}{\partial x_j} + \frac{\partial \widetilde{u}_j}{\partial x_i}\right), \tag{4.6}$$

where $C_S$ is the Smagorinsky constant, which is corrected near the wall following [109], as

$$C_S = \left(C_0^n + \left\{\kappa\left(\frac{z}{\Delta} + \frac{z_0}{\Delta}\right)\right\}^{-n}\right)^{-1/n}, \tag{4.7}$$

where $z_0$ is the roughness length scale, $\Delta = \sqrt[3]{\Delta x \Delta y \Delta z}$ the grid size and choosing $(C_0, n) = (0.14, 3)$. Finally, $-\partial_i p_\infty$ is a constant pressure gradient applied to flow to reproduce the effect of geostrophic forcing. It is worth noting that in equation (4.4) the viscous term of the Navier-Stokes equation has been neglected. This can be done when the Reynolds number is high enough, added to the fact that the mesh is coarse enough so

that the near-ground viscous layer will not have any impact on the calculations. Another implication of being under-resolved in the near-ground surface is that applying a no-slip boundary condition is devoid of any meaning. Instead, we simulate the bulk effect of velocity shear by replacing it with slip velocity and imposing a local shear stress through a wall-stress model. Commonly, the wall-stress model of [119] is used by considering the Monin-Obukhov similarity theory (MOST) theory and in this study we employ the recent formulations of [15],

$$\tau^{\text{wall}}(x,y) = \tau_w(x,y)\frac{\widehat{\widetilde{u}}_i(x,y,\Delta z/2)}{\sqrt{\widehat{\widetilde{u}}_x^2(x,y,\Delta z/2)+\widehat{\widetilde{u}}_y^2(x,y,\Delta z/2)}}, \tag{4.8}$$

and

$$\tau_w(x,y) = -\left[\frac{\kappa}{\ln\left(\frac{\Delta z/2}{z_0}\right)}\right]^2 \left[\widehat{\widetilde{u}}_x^2(x,y,\Delta z/2)+\widehat{\widetilde{u}}_y^2(x,y,\Delta z/2)\right], \tag{4.9}$$

in which the horizontal-averaged boundary condition is replaced by a twice-filtered velocity. In addition, we make use of periodic boundary conditions in lateral directions and free/partial-slip ( $\partial_3\widetilde{u}_i = \widetilde{u}_3 = 0$ for $i \in 1,2$) for the bottom $z = 0$, and the top boundary $z = H$, respectively. To solve the governing equations, sixth-order compact finite-difference schemes are used [97] on a Cartesian mesh in a half-staggered arrangement (the same mesh is used for the three velocity components $(u_1, u_2, u_3) = (u, v, w)$, with a different mesh used for pressure $p$), pressure correction is achieved by a direct spectral solver, while for the time integration a low-storage third-order Runge-Kutta scheme is employed. Additionally, explicit filtering is also applied to all fields and at each time step using discrete sixth-order filters [58]. Finally, massive parallelisation of the numerical solver is achieved using MPI and an efficient 2D pencil domain decomposition approach [98]. The current implementation has been run on over $\mathcal{O}(10^5)$ number of processors and in multiple platforms and has consistently shown excellent scaling properties.

## 4.2 Comparing turbulence generation methods

For a robust comparison between the different synthetic turbulence models (von Kármán, Kaimal, and LES), we must first ensure that they all yield statistically similar results. This includes both the time-averaged velocity profiles as well as velocity spectra. Here, the neutral atmospheric boundary layer on a flat Earth is considered with three levels of

surface roughness, $z_0 = 0.0001\,\mathrm{m}$, $0.001\,\mathrm{m}$ and $0.01\,\mathrm{m}$. This defines three flow conditions that will be referred to as *Mild*, *Regular* and *Severe*, and correspond to surface friction velocities $u_* = 0.174\,\mathrm{m\,s^{-1}}$, $0.235\,\mathrm{m\,s^{-1}}$ and $0.274\,\mathrm{m\,s^{-1}}$, respectively. The friction velocities have been selected such that the mean velocity at altitude 10 m attains a value of approximately $U_{10} \approx 5\,\mathrm{m\,s^{-1}}$ in all cases. Further details are provided in the table 4.1 below. Before proceeding to present the results it is also instructive to report the simula-

**Table 4.1** Atmospheric turbulence flow parameters

| Name | $z_0$ [m] | $u_*$ [m/s] | $U_{10}$ [m/s] | $\mathcal{I}$ % |
|---|---|---|---|---|
| *Mild* | 0.0001 | 0.174 | 5.05 | 8.87 |
| *Regular* | 0.001 | 0.235 | 5.41 | 10.74 |
| *Severe* | 0.01 | 0.274 | 4.74 | 12.73 |

tion configurations and some key parameters used for generating the data. Starting with large-eddy simulations, we consider a region of $L_x \times L_y \times L_z = 1000\,\mathrm{m} \times 500\,\mathrm{m} \times 500\,\mathrm{m}$ discretised in $N_x \times N_y \times N_z = 1024 \times 512 \times 513$ mesh nodes. The time step was selected to be equal to $\Delta t = 0.05\,\mathrm{s}$ so that a Courant-Friedrichs-Lewy (CFL) number of approximately 0.25 is maintained throughout. Each simulation was run for 50 large-eddy turnover times ($\delta/u_*$, where $\delta = 1000\,\mathrm{m}$ is the boundary layer height) a spin-up time that has been previously used to obtain converged statistics [15, 210]. Restricted by the CFL number, the total number of time steps required during the precursor simulation amounts to 2 million. All three cases were run using 4096 processors and required approximately 1.5 million CPU hours each. For the simulations involving the aeroelastic coupling, a small number of snapshots were extracted (every 3 seconds) to obtain wind fields for a total of 10 minutes. For the Kaimal model a smaller mesh of $N_y \times N_z = 32 \times 32$ nodes was used for a domain of $L_y \times L_z = 50\,\mathrm{m} \times 50\,\mathrm{m}$, obtaining a resolution of approximately 1.5 m. The calculations were made for a total analysis time of 250 s using a time step $\Delta t = 0.25\,\mathrm{s}$. To reproduce the desired results TurbSim [87] was used in a customised mode by applying a logarithmic profile and setting the turbulence level at the characteristic altitude, surface frictions and coherence parameters so that the LES results are closely matched. Lastly, for the von Kármán model, the standard deviations of the three velocity components $\sigma_u$, $\sigma_v$ and $\sigma_w$ together with the length scale $L_u$ extracted from the LES simulations were used to generate zero-mean time series, which were then added to the background characteristic velocity $U_{10}$. A comparison of the vertical profiles and power spectral densities (PSD) are shown in figures 4.1 and 4.2, respectively. The vertical mean velocity profiles agree well for the Kaimal and LES data whereas the von Kármán

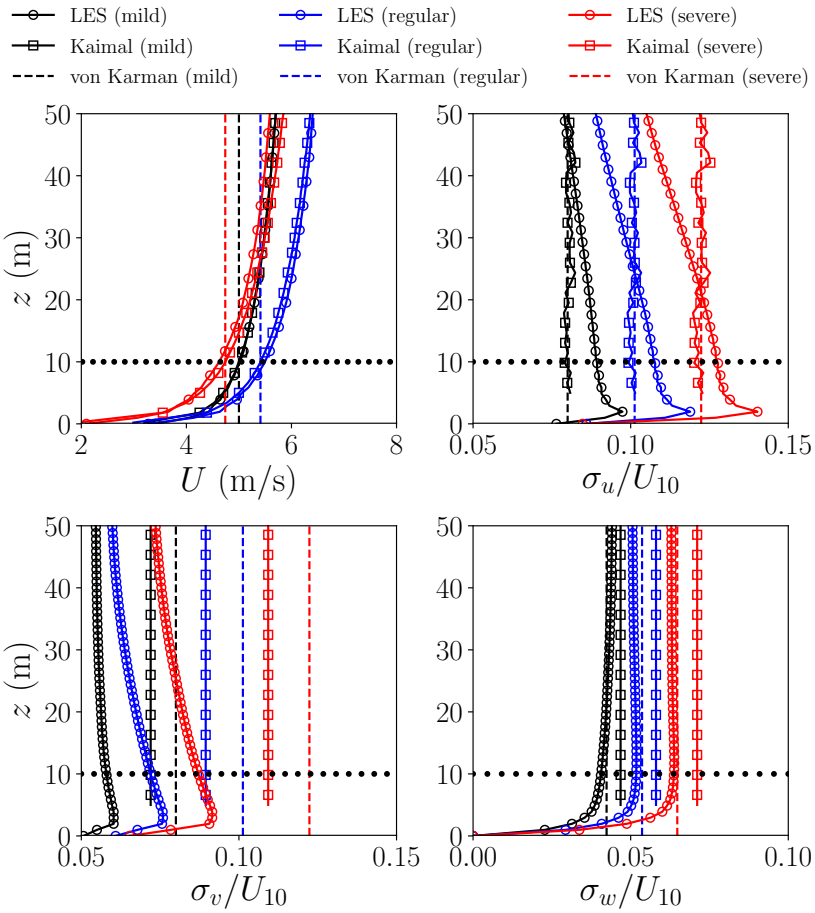



**Fig. 4.1** Mean velocity and turbulence intensity profiles for the three cases. The dotted lines show the reference altitude $H_{\text{ref}} = 10\,\text{m}$.

can only match the velocity at the reference altitude (10 m). Larger discrepancies can be found between the three models in the streamwise and transverse velocity variances. In particular, LES predicts a linear decrease of the streamwise variance, a result that agrees with the observed data as well as theoretical scaling laws [208], whereas the Kaimal and von Kármán models yield uniform distributions for all components. Similarly, PSDs obtained at the reference height $H_{\text{ref}} = 10\,\text{m}$ from the three models, shown in figure 4.2, appear to collapse together for the streamwise component but exhibit larger discrepancies in the lateral ones. Specifically, in the vertical velocity component PSDs, LES data plots significantly deviate from those of Kaimal and von Kármán. Those two models assume flow isotropy which results in nearly identical PSDs over the $v'$ and $w'$ components. Nevertheless, while our comparison of the flow characteristics is ended here by comparing at most second-order statistics (velocity variance, spectra) more differences between the LES,

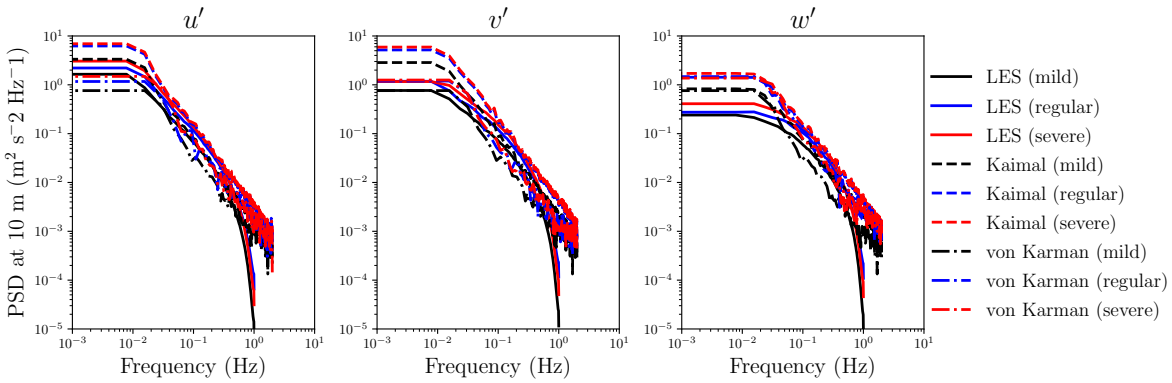



**Fig. 4.2** Power spectral density of velocity fluctuations for each velocity component ($u'$, $v'$ and $w'$) and all three models (1D von Kármán, 2D Kaimal and 3D LES).

which produces a non-Gaussian distribution of velocity, and the stochastically-generated turbulence can be further observed. This is for example demonstrated qualitatively in figure 4.3. The two snapshots of the LES and Kaimal, respectively, are both taken from the *Mild* case and exhibit large differences in the formation of the near-wall large coherent structures. LES flow structures consist of large counter-rotating stream tubes generating large structures which extend as high as 25 m while Kaimal consists of smaller structures extending only a few metres along the $z$-direction. The von Kármán model on the other hand is not plotted as only a uniform distribution would be seen for this kind of slice. These qualitatively observed differences between the three models which are not captured by the (up to second order) flow statistics will be re-assessed later on during the aeroelastic simulations in Section 4.4.

## 4.3 3D velocity fields as aeroelastic inputs

A final aspect of the present model is the methodology for combining the turbulence input data with the dynamic aeroelastic solver. Atmospheric turbulence fields either in a two-dimensional form (vertical planes) or three-dimensional sub-domains of a flow realisation where generated for a time period of 120 s. Subsequently, the data was stored and read by SHARPy at each time step. The LES input data was found to impose practical constraints in the reading and interpolation process as each snapshot amounts to more than 3 Gb of binary memory (for $1024 \times 512 \times 513 \sim 270$ million grid points). This poses a bottleneck which may be approached in two ways. A direct solution is to read and store in RAM only the two necessary snapshots. This is suitable for distributed

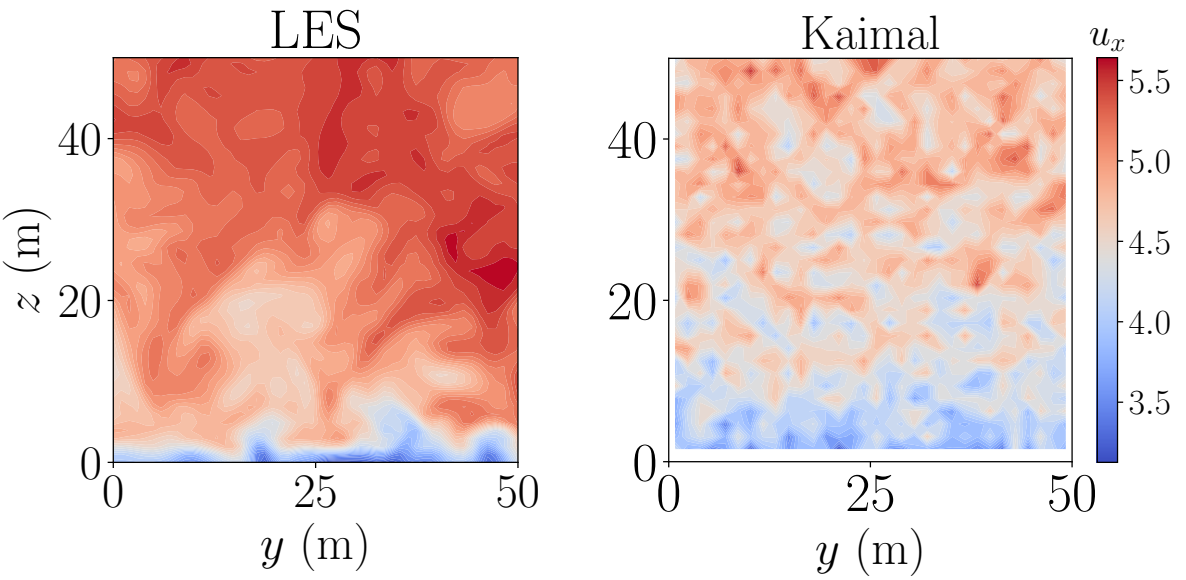



**Fig. 4.3** Instantaneous snapshots of a *mild* streamwise velocity, $u(t)$, for the Kaimal and LES models.

computers, where permanent storage latency is high, and nodes can be devoted to a small number of simulations each. However, shared memory computers usually have a lower permanent storage latency and the RAM available per processor node is lower. The solution is then to map the binary file stored in a relatively fast hard drive into an array-like structure, so data is directly read from the permanent storage. This solution was seen to drastically reduce the memory overhead with only a 20% increase of the interpolation time cost. The 3D velocity field reader implemented here supports both methods, and it implements a cache where data fields are read only once during each simulation. In the end, the decision of which method is used for the 3D field read is a combination of the architecture the software is being run in, the number of cases running simultaneously and the time between LES snapshots. Once the 3D field is mapped or copied into an array, a multivariable interpolation is conducted. First, the volumetric data is given as a structured regular grid. This makes trilinear interpolation a suitable solution, while the quicker closest neighbour interpolation is also available for finer 3D field discretisations. Results presented here are obtained using a trilinear interpolation in space. The time-dependency of the 3D velocity fields is considered in this analysis. When multiple snapshots of the field are given, a linear interpolation is performed between the two closest snapshots to the simulation time. A visualisation of the combined 3D velocity fields with an overlay of a typical aircraft trajectory is shown in figure 4.4 below.

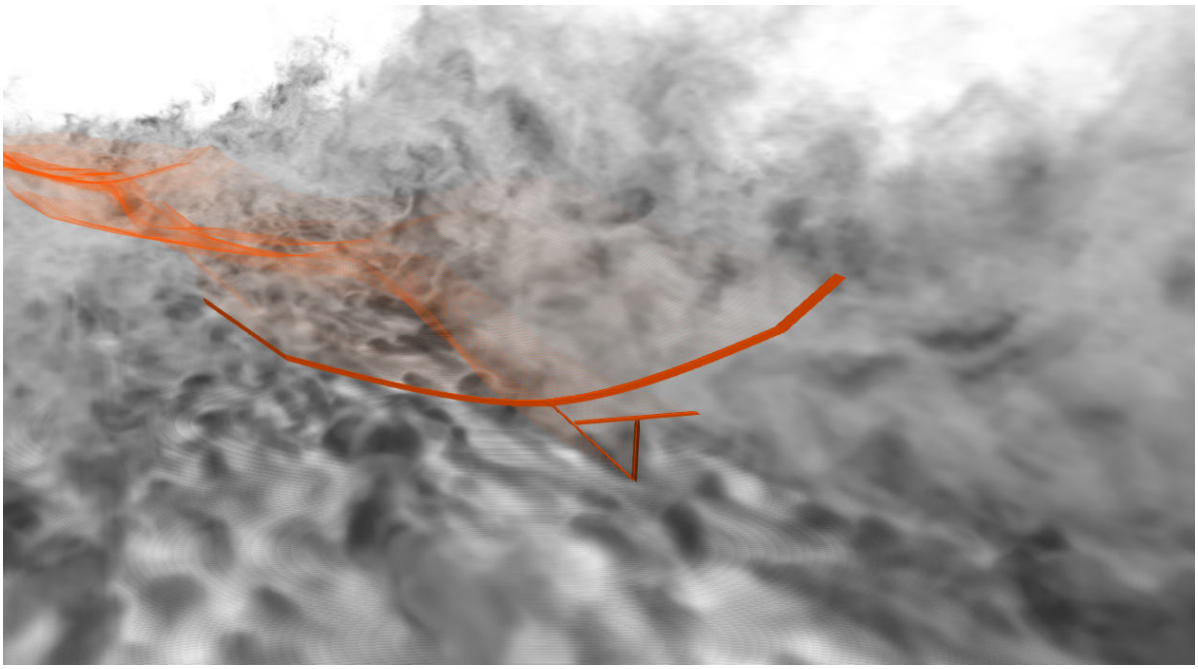


**Fig. 4.4** Snapshot of the HALE aircraft inmersed in the severe turbulence 3D LES field. The aerodynamic and structural grids are shown in orange, together with their time history for the next 30 seconds. The turbulent field is shown as a mapping of horizontal velocity, where low velocities are opaque and high velocities are increasingly transparent.

## 4.4 Numerical results

As discussed in Section 4.1, the 1D von Kármán spectrum is typically used to simulate continuous turbulence in aircraft structural design. Having already presented the differences between the three turbulence-generation models (1D von Kármán, 2D Kaimal and 3D LES), in this section we shall attempt to quantify their effects on aircraft aeroelasticity from short simulated flights in open loop (that is, without feedback control on the vehicle dynamics). All simulations in this work start with a trimmed vehicle at the reference height $H_{\text{ref}} = 10\,\text{m}$.

### 4.4.1 Flexible aircraft dynamics for each atmospheric turbulence model

To examine the relative importance of input turbulence data in predicting both loads and flight trajectories for a flexible aircraft we undertake a number of simulations for the T-tail HALE model (Section 3.6) in all three turbulence intensity (TI) cases. Three-dimensional atmospheric turbulence fields are extracted from the three models such that they exhibit

**Table 4.2** Summary of the simulations used in this chapter calculations, showing the similarity in statistics between the von Kármán, Kaimal and LES fields.

| Case | $U_{10}$ (m/s) | $\mathcal{I}_{10}$ (%) | Y($t = 0$) (m) | Turbulence model |
|---|---|---|---|---|
| L11 | 5.05 | 8.87 | 0 | LES |
| L12 | 5.05 | 8.87 | 10 | LES |
| L13 | 5.05 | 8.87 | 50 | LES |
| L14 | 5.05 | 8.87 | 200 | LES |
| L15 | 5.05 | 8.87 | -200 | LES |
| L21 | 5.41 | 10.74 | 0 | LES |
| L22 | 5.41 | 10.74 | 10 | LES |
| L23 | 5.41 | 10.74 | 50 | LES |
| L24 | 5.41 | 10.74 | 200 | LES |
| L25 | 5.41 | 10.74 | -200 | LES |
| L31 | 4.74 | 12.73 | 0 | LES |
| L32 | 4.74 | 12.73 | 10 | LES |
| L33 | 4.74 | 12.73 | 50 | LES |
| L34 | 4.74 | 12.73 | 200 | LES |
| L35 | 4.74 | 12.73 | -200 | LES |
| K1 | 5.05 | 8.87 | - | Kaimal |
| K2 | 5.41 | 10.74 | - | Kaimal |
| K3 | 4.74 | 12.73 | - | Kaimal |
| VK1 | 5.05 | 8.87 | - | von Kármán |
| VK2 | 5.41 | 10.74 | - | von Kármán |
| VK3 | 4.74 | 12.73 | - | von Kármán |

similar flow statistics, as discussed in section 4.2, with simulations corresponding to 120 s of real time and $\Delta t = 0.025$ s. The produced time-series for the aircraft loads and trajectories for each TI and turbulence-generation model are subsequently used in our assessments. For the LES-generated background flows, multiple aeroelastic simulations were run for each turbulent intensity case in order to explore the spanwise flow variations of the simulated atmospheric turbulence. A summary of all aeroelastic simulations is shown in table 4.2 below. Each case is defined by the average reference velocity and turbulence intensity as well as the initial spanwise position in the turbulent flow field, $Y$, of the aircraft mid-wing point. Values of $Y = 0$ m, 10 m, 50 m, 200 m and $-200$ m are selected based on the lateral turbulence integral length scale, which is calculated to be between 150 m to 200 m for the three turbulence intensity cases.

Aircraft trajectories are tracked at the origin of the vehicle axes shown in Figure 2.1 and the trajectories for all simulations in Table 4.2 are shown in Figure 4.5. We may observe that the aircraft exhibits a similar vertical trajectory for all turbulence

inputs by maintaining an altitude around 10 m with "dives" and "jumps" of around 1 m to 3 m occurring throughout. A significant deviation from this pattern is observed in the response to the von Kármán model for both the regular and severe cases. Perfect spanwise coherence in the background flow in this case implies generation of a net resultant lift that pushes the vehicle upwards, which can most clearly be seeing under "regular" turbulence level, although the same effect on the elevator may make the aircraft to veer downwards, as it seemingly occurs in the "severe" turbulence level. As it will be see later, this is associated to a reduction in wing loads. The transverse component of the trajectory exhibits larger deviations between the three turbulence-generation models as well as within LES itself when the aircraft is initially placed at different spanwise locations with respect to the same background flow ($Y = 0$ m, 10 m, 50 m, 200 m and $-200$ m). Interactions with the large scale flow features in the LES-generated turbulence result in significant variation of the aircraft trajectory, which can deviate both right and left while in some cases it may experience a full U-turn and continue flying in the opposite direction. The other two models (Kaimal and von Kármán), which are characterised by a smaller and no variability in the transverse direction, respectively, result in nearly straight flight trajectories. The effect of spanwise velocity coherence experienced by the aircraft wing, as well as its relationship with the aerodynamic behaviour (e.g. spanwise lift distribution) and aircraft trim shape, is investigated in figure 4.6. It shows the trajectory of the aircraft initiated at $Y = 50$ m spanwise distance from the reference point at the centre of the LES computational domain, under LES-generated inflow with "regular" TI level. The figure shows the trajectory of the aircraft on which three different segments, of 25 s each, have been highlighted. They are referred to as A, B, and C and the corresponding ensembled-averaged spanwise velocity and lift distributions along the aircraft wing are shown. Clearly, as the aircraft comes across large coherent structures it experiences an incident velocity with large variations across its span. In turn, the non-uniform velocity distribution gives rise to asymmetric loads which through the aircraft flexibility lead to a gradual veer to the right, which finally leads to a full U-turn. It should be noted that the aircraft altitude does not change significantly during the U-turn trajectory and that ensemble averaged lift distributions reduced further this variability. This is a result of the inherent stability that most aircraft present. This of course can be modified by a flight control system which can alter the aircraft trim and thus have an impact on loads as well. As far as the root loads are concerned, similar variability is observed for all three atmospheric turbulence inputs. In figure 4.7 we observe that torsion fluctuates around a mean value of 3.5 N m for all three models with larger amplitudes found for the regular

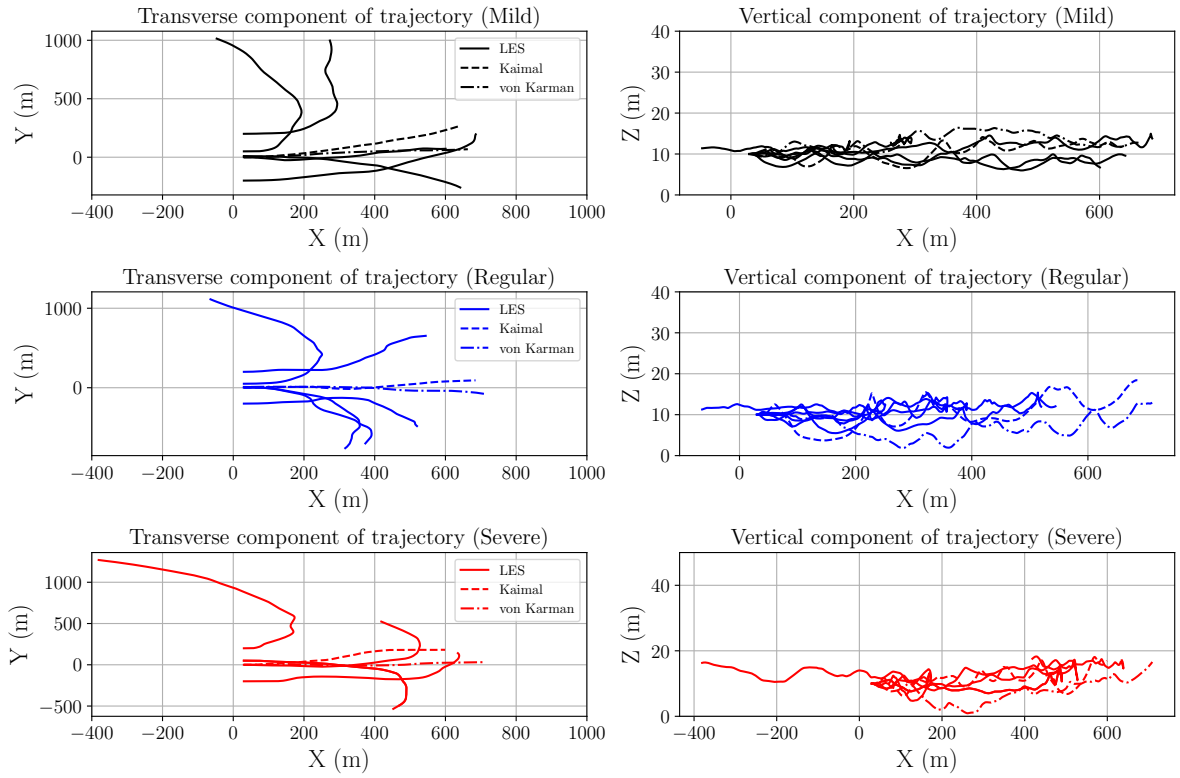



**Fig. 4.5** Sample flight trajectory of the flexible configuration and all three turbulence generation methods

and severe Kaimal cases. Out-of-plane (OOP) and in-plane (IP) wing root moments, on the other hand, attain a mean value of around −1700 Nm and 370 Nm, respectively, with the fluctuations under the Kaimal and LES inputs being more pronounced in the time series. It is worth mentioning here that the aeroelastic simulations which make use of the Kaimal model exhibit relatively high-frequency fluctuations. These fluctuations which appear to dominate the loads after about 12.5 s coincide with the instance that the aircraft trajectory deviates from a straight path (the one aligned with the mean flow). This effect can be attributed to the lack of streamwise coherence of the Kaimal model as well as the assumption of "frozen turbulence" which generates non-physical longitudinal structures once the 2D planes are projected to a three-dimensional wind field. Once the aircraft is misaligned from the mean flow, it encounters a non-uniform velocity field with large velocity differences across its wing span, increasing thus the torsion experienced by the root. Similar fluctuations can also be observed for the "severe" case and the LES inputs. This however is not an artefact of the LES method but rather the effect of the near-surface large coherent flow structures. Nonetheless, in order

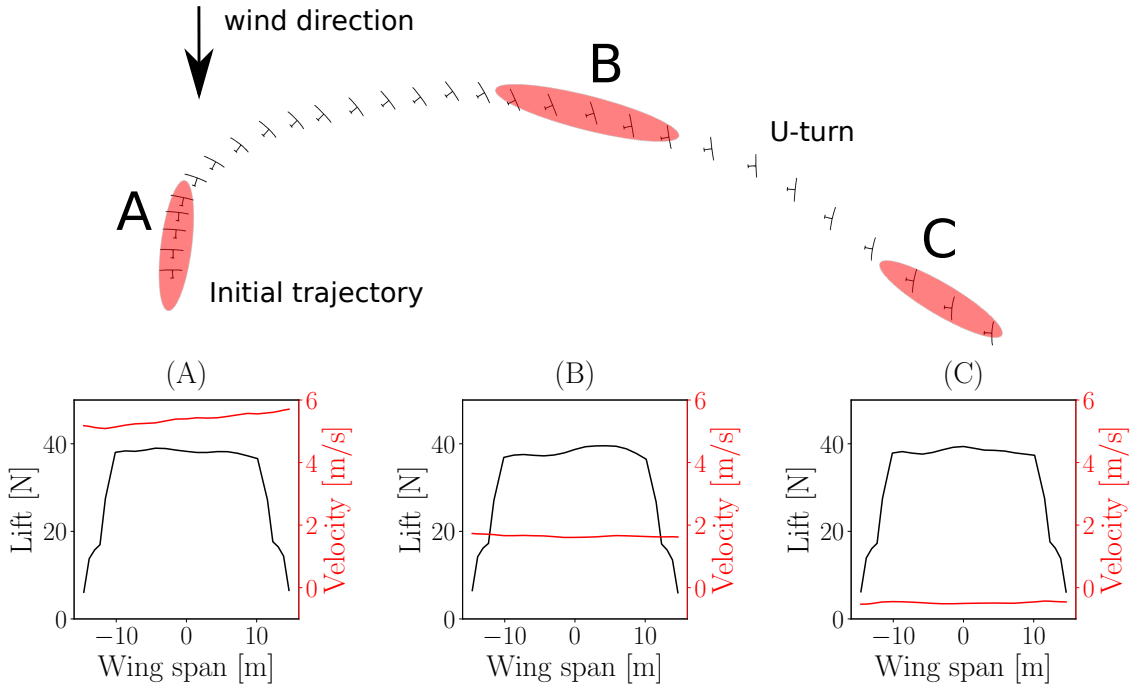



**Fig. 4.6** Sample flight trajectory and ensemble-averaged lift/velocity distribution of the flexible configuration for the regular LES case

to fully appreciate the magnitude of the loads it is worth plotting the load envelops for all three TI levels and root loads of interest, namely torsion, OOP and IP. The load envelops which are constructed by the convex hull obtained from the maximum instantaneous loads exerted on the aircraft, provide information regarding the loading regime expected for each TI case. To this end, from both figures 4.8 and 4.9 we may observe the gradual expansion of the load envelopes as the input turbulence intensity increases. The shape of the load envelopes is also consistent in all three turbulence intensity cases. In the case of the OOP plotted against torsion, the load envelopes are stretched along the OOP direction signifying the significant increase of OOP with increasing levels of turbulence. Similarly, the shape of the IP versus torsion envelopes appears to be elongated with the larger in-plane moments occurring at negative root torsion. In both cases (OOP and IP) the choice of the turbulence input model appears to be more significant. The LES data provide by far the largest load envelope with its boundaries twice as big as those of Kaimal. On the other hand, the Kaimal model yields smaller load envelopes with the difference between the two models being increased with

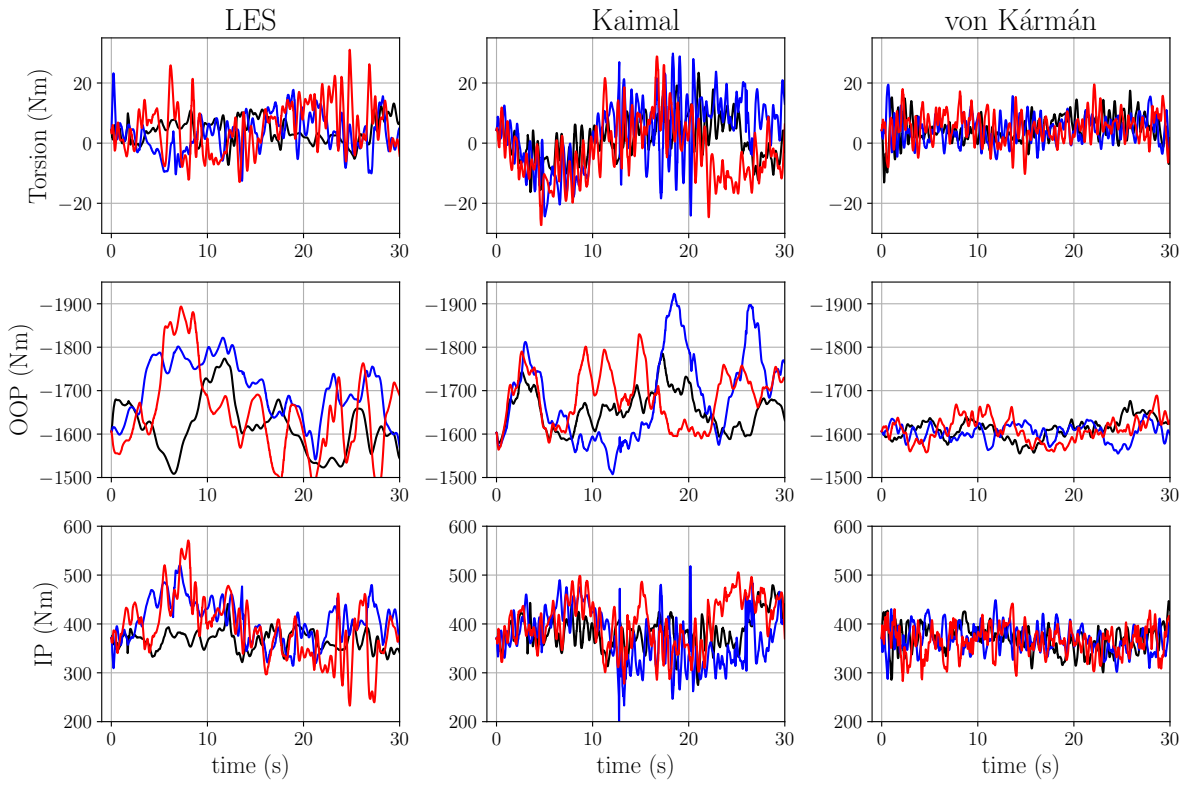



**Fig. 4.7** Root loads from the three turbulence generation models. In this figure, torsion, out-of-plane (OOP) and in-plane (IP) moments are plotted as a function of time for the mild (black), regular (blue) and severe (red) turbulence levels.

the magnitude of turbulence intensity. Perhaps surprisingly, the smallest load envelopes are formed by the von Kármán model input. This is true for both the IP and OOP moments. We argue that this difference is due to the fact that in the von Kármán model the perfect spanwise coherence gives uniform spanwise load distribution. As a result, the dynamic loads give rise to a lifting of the overall structure through "inertial relief" (applied work by gust forces is transformed into rigid body kinetic energy, instead of structural deformation), which is observed in a higher spectral content of vehicle accelerations, as will be seen later, and results in the larger vertical excursions seen in figure 4.5 for all three turbulence intensity cases of the von Kármán model. Another important difference that is captured by the load envelopes is the frequency of occurrence of critical loading events. We may observe that the boundaries of all Kaimal and LES load envelopes are due to intermittent wind gusts which increase with increasing levels of turbulence. To determine the potential effects of intermittency on the aircraft loads, we have computed the probability distribution function (PDF) of the incident velocity and

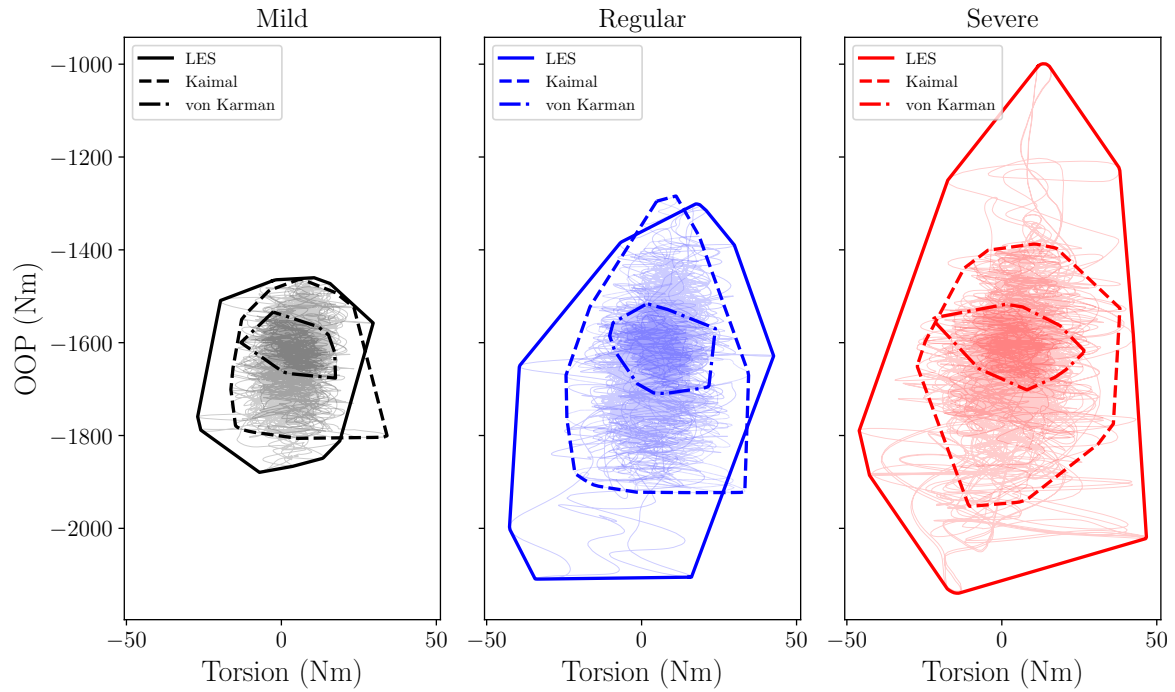



**Fig. 4.8** Root loads envelopes for all three models and cases. In all figures the out-of-plane (OOP) root moment is plotted against torsion.

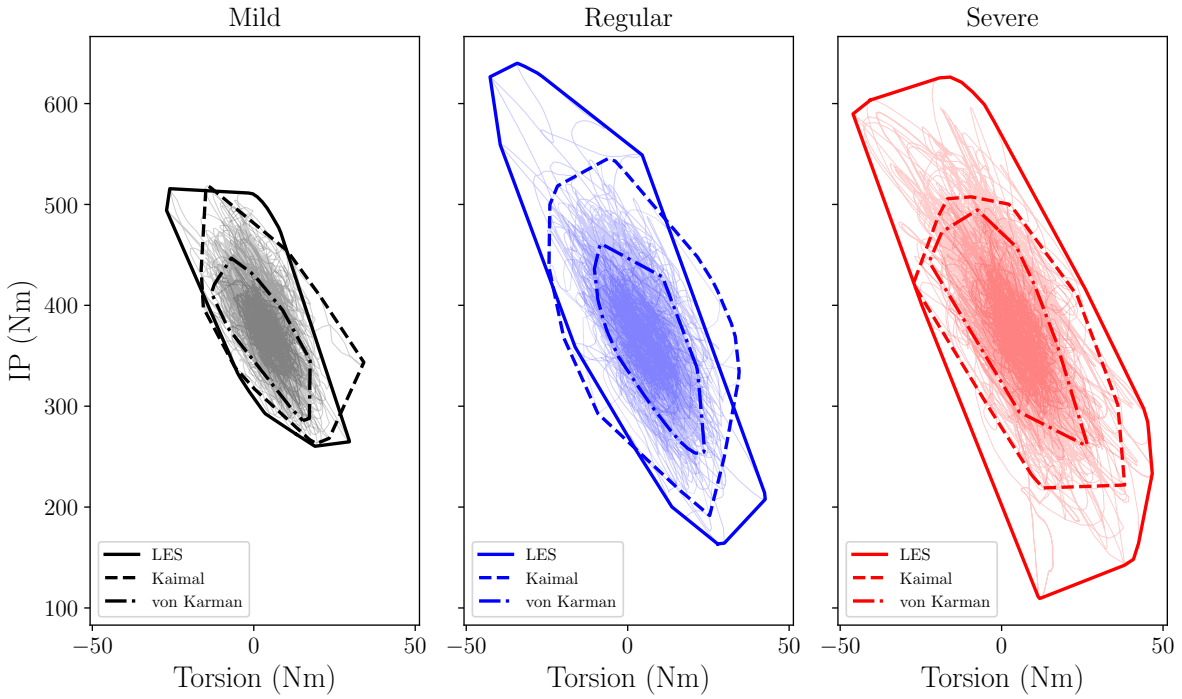



**Fig. 4.9** Root loads envelopes for all three models obtained from combining all available time histories. In all figures the in-plane (IP) root moment is plotted against torsion.

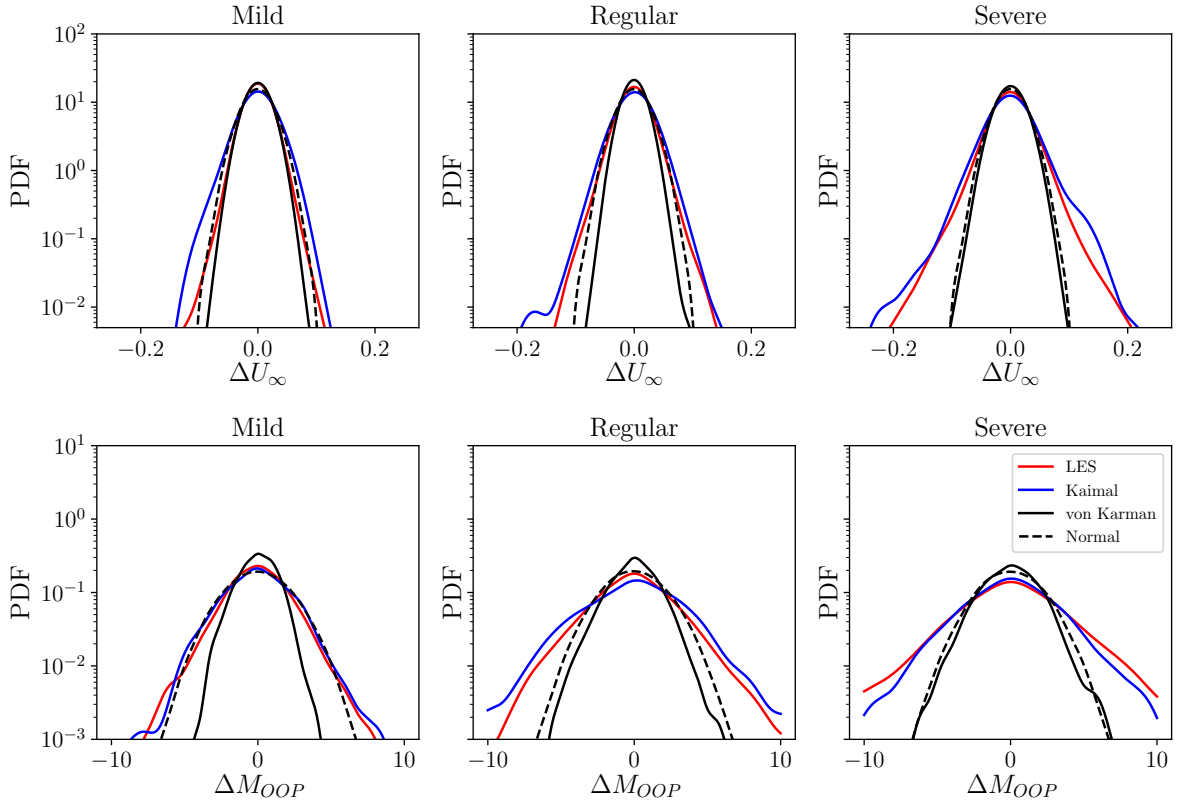



**Fig. 4.10** PDF of the incident velocity magnitude$\Delta|U_\infty|$ and out-of-plane moment $M_{OOP}$ for all models and turbulence intensity cases. Normal distribution is included for comparison.

load increments ($\Delta u = u(t+\Delta t) - u(t)$, $\Delta M = M(t+\Delta t) - M(t)$) and examined the characteristics of its tails. Intermittency in both the incident velocity and the loads can be captured by heavy tail distributions that significantly deviate form the standard Gaussian distribution. Figure 4.10 illustrates the existence of both intermittent velocity difference and loads, particularly for the "severe" case in which a clear heavy tail distributions, $|\Delta U_\infty| > 0.1\,\mathrm{m\,s^{-1}}$ and $|\Delta M_{OOP}| > 7.5\,\mathrm{N\,m}$ in both at positive and negative. It is worth noting that in the incident velocity difference PDFs the Kaimal model appears to be more intermittent. This is again due to the frozen turbulence assumption in combination with the deviation of the aircraft from a straight path. The wind field intermittency also translates into the observed loads. More specifically, the PDFs for $|\Delta M_{OOP}|$ show that the aircraft experiences large variation of out-of-plane loads sporadically with probability density of less that 0.01, particularly in the "severe" case. In the other two "milder" cases the two models (Kaimal and LES) appear to be in a better agreement. Von Kármán turbulence is characterised by significantly less variability and thus, its PDF appears much narrower. Comparing the OOP moment distribution with the other two models as

well as with that of the incident velocity for the von Kármán model we may not be able to infer the origin of such a large difference. However, following our earlier argument we may again justify such a large difference in the PDF occurs by the stronger inertia relief in this case. Finally, it is worth introducing the power spectral density (PSD) functions of the root loads as well as vertical acceleration of the wing root ($a_z$) from all the three models in figure 4.11. The comparison of PSDs for the three models reveals a similar behaviour between the different levels of turbulence. For instance, the PSD of torsion is characterised by a plateau region over the low-frequency region, followed by a steep slope and low-energy content ∼0.001 in the high frequencies. In addition, all three turbulence intensity cases exhibit a peak around $f = 1.3\,\text{Hz}$ which corresponds to the mode II (natural torsion frequency). The peak around mode II is more pronounced for the severe turbulence case, while modes I ($f = 0.4\,\text{Hz}$) and V ($f = 3.92\,\text{Hz}$) are highly damped (figure 3.23) and are not discernible in the PSD plots. Nonetheless, the effect of increasing the level of turbulence intensity on each PSD can be seen in all three cases by "pushing" the distributions up. This effect is clearer over the lower frequencies (0.1 Hz to 1 Hz) and in particular for the simulations using the LES-generated turbulence input. A similar trend is followed by the IP moment. On the other hand, the OOP follows a smoother slope and lacks a "plateau" region, confirming that its response varies across the range of scales. The PSDs for the load factor, $N_z = a_z/g$, in figure 4.11 show that the response to von Kármán turbulence is consistently above the LES and Kaimal ones at the lower end of the frequency spectrum, which is the one determining the flight dynamics of the aircraft. There is therefore a clear exchange between lower loads and higher accelerations (i.e., inertia relief) in the response to the von Kármán turbulence model, which is ultimately responsible for the lower load levels found in figures 4.8 and 4.9.

### 4.4.2 Assessment of aeroelastic effects

Having demonstrated the differences between three turbulence input models (von Kármán, Kaimal and LES), in this section we present a cross-comparison between the *Flexible* and the *Stiff* aircraft configurations. For this analysis we only use the LES turbulence inputs. The two aircraft configurations are expected to exhibit a different response to dynamic gust loading. The flight dynamics of the former (*Flexible*) aircraft configuration are characterised by its coupling with the low-frequency structural modes [2]. On the other hand, the *Stiff* aircraft configuration is expected to exhibit a substantial degree

of decoupling between aeroelasticity and flight dynamics. To address the question of how much does stiffness affect the flight dynamics and structural loads of an aircraft subject to atmospheric turbulence, we present and compare the load envelopes and PSD distributions of the two configurations. To this end, the two aircraft configurations have been simulated for the same number of cases by maintaining all other parameters the same. We begin with the presentation of the OOP vs torsion load envelope in figure 4.12. Aircraft flexibility is found to fundamentally change the dynamic gust response of the aircraft. The stiff aircraft is found to be prone to larger OOP root moment loads and suffer less from large variations in torsion. More specifically, OOP moment attains values ranging from $-500\,\mathrm{N\,m}$ to $-2500\,\mathrm{N\,m}$ while experiencing only positive torsion values throughout. This shift in the loads response is also reflected in the respective PSDs of the loads. Figure 4.12 shows a comparison between the two cases for all computed root loads. As in the load envelopes, torsion fluctuations of the stiff case are an order-of-magnitude smaller than those of the flexible one with the differences being more pronounced in the low-frequency region. Conversely, the OOP PSD shows a significant increase in the energy over the mid-range frequencies (inertial subrange). Finally smaller difference can be observed for the IP moment with the stiff case exhibiting smaller fluctuations. This transfer of energy from torsion and IP moment fluctuations to OOP signifies the fundamental differences in the turbulence/loads transfer function between the flexible and the stiff aircraft. The former transforms the OOP loads into torsion which in turn changes its aerodynamic behaviour and therefore the resultant root loads. This can be seen as a passive load alleviation mechanism which targets a reduction of the OOP moment. The stiff aircraft on the other hand, by experiencing a certain degree of decoupling between structural and flight dynamics, remains more responsive to incident turbulence (the $-5/3$ scaling is maintained in the inertial sub-range) implying it may experience larger fatigue loads and ride quality during its lifetime.

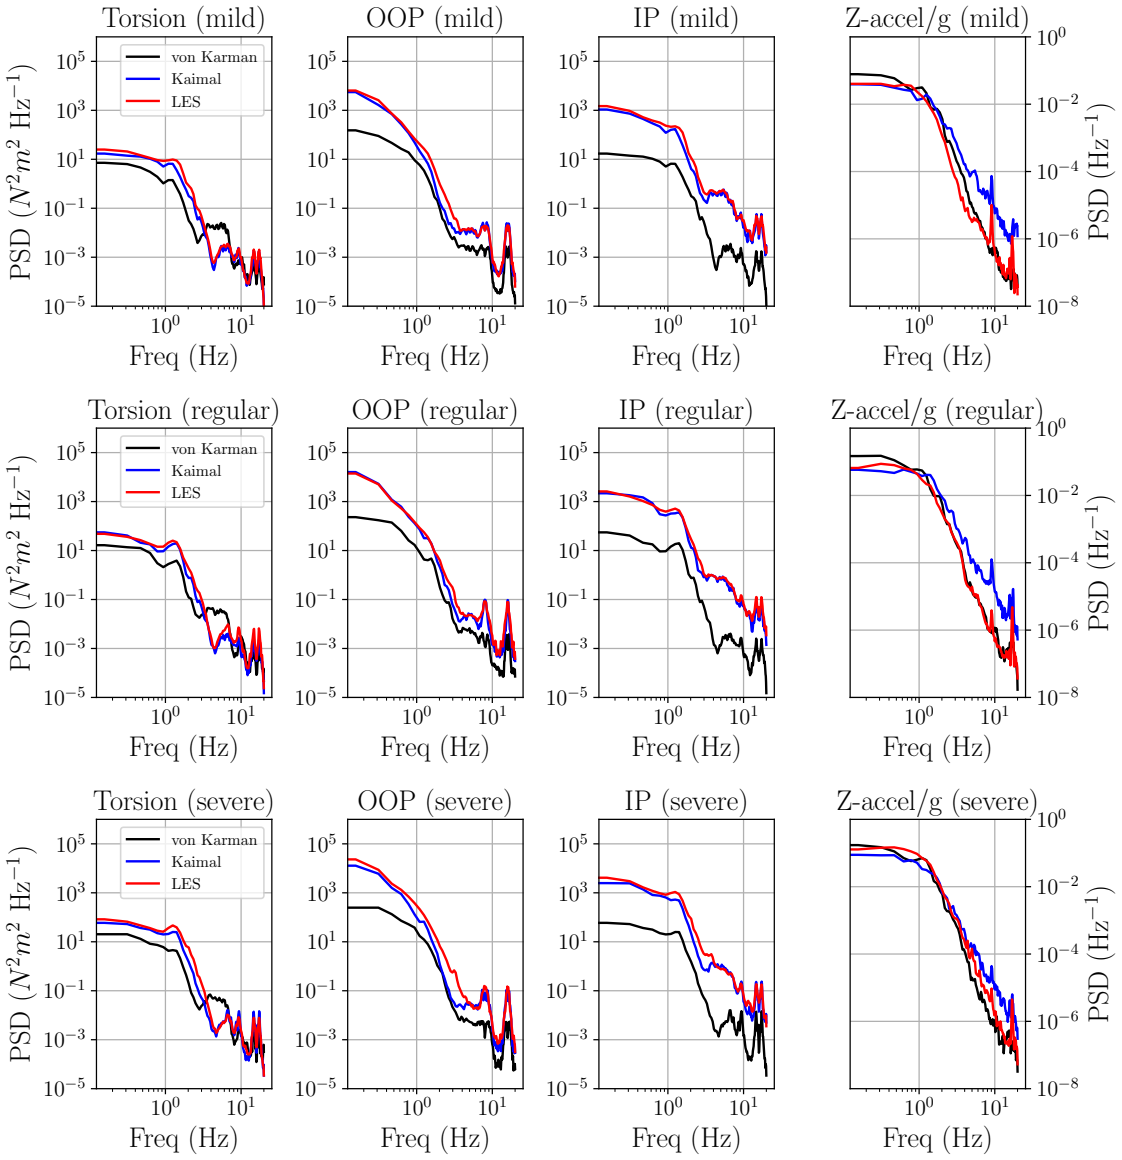



**Fig. 4.11** PSDs of the root loads (Torsion, OOP and IP) and normalised vertical acceleration at the root for all three turbulence input cases (LES, Kaimal and von Kármán) and turbulence inputs (mild, regular, severe).

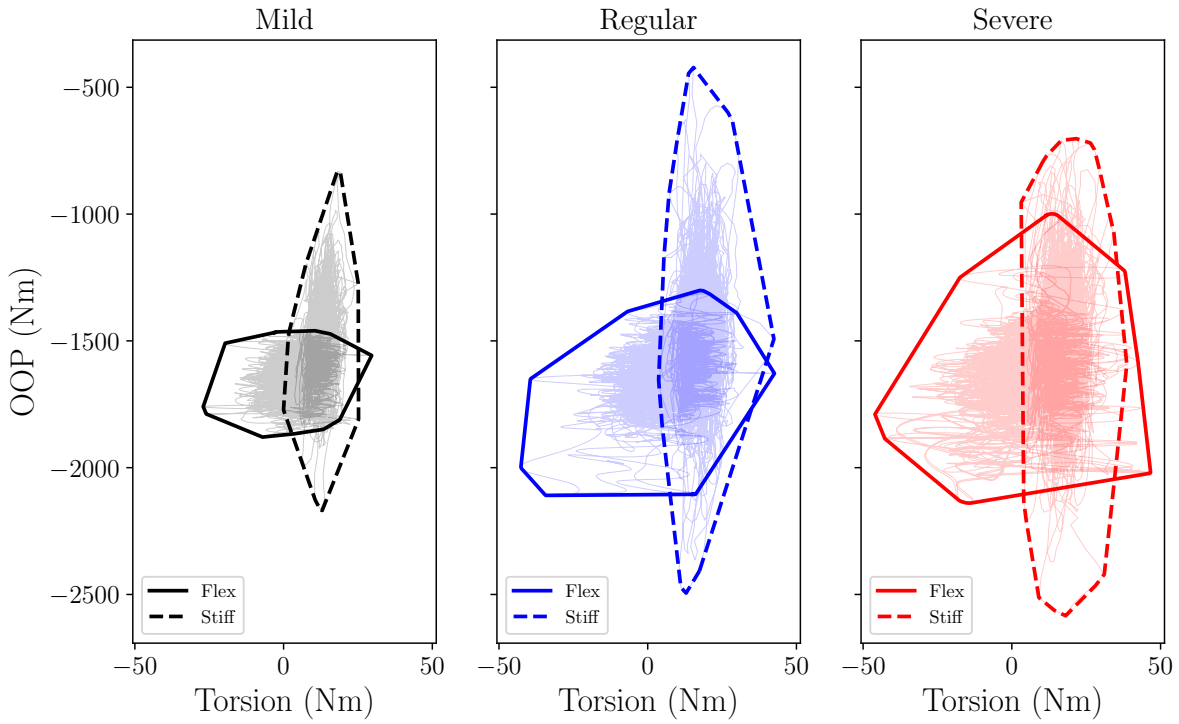



**Fig. 4.12** Root loads envelopes for the "Stiff" and "Flexible" cases for all levels of TI. In all figures the out-of-plane (OOP) root moment is plotted against torsion.

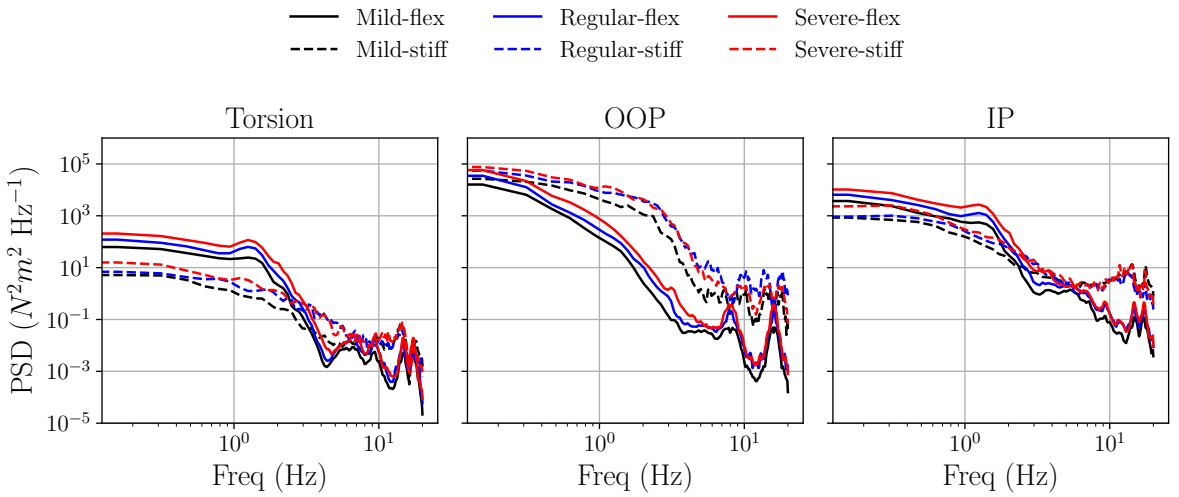



**Fig. 4.13** Comparison of PSD distributions for the root loads (Torsion, OOP and IP) between the stiff and flexible aircraft.

# Chapter 5

# A complex nonlinear manoeuvre optimisation problem

A second problem of interest at low altitudes relates to the safe deployment of the vehicle. This chapter will describe the modelling and implementation details, and the parameter choices for simulating and optimising a complex, non-linear manoeuvre. As a working example of real interest, the Catapult-Assisted Takeoff (CATO) of a representative solar-powered HALE aircraft has been chosen. Robust and scalable launching procedures are of special interest as HALE vehicles evolve from experimental to commercial products. Depending on the total mass and geometry, the launch procedures found in the different prototypes under development include manual launch, a conventional runway takeoff or towing on a trailer. In particular, this procedure offers reliability, predictability and robustness, while the relatively small footprint of the catapult would allow to easily modify the takeoff orientation. Importantly, takeoffs are good examples of aeroelastic simulation where no linearisation can be carried out. The instantaneous aircraft velocity evolves from rest to close to cruise speed, while the structural deformation varies from a multiply-supported gravity-only loading condition to an aeroelastically trimmed shape. This happens in a few seconds. Finally, slow wing bending frequencies imply that, for short enough ramps, the state of the aircraft on release might be highly dynamic and far from a quasi-steady climb condition.

The efficiency of SHARPy allows to optimise this complex manoeuvre given a parametrisation of the ramp geometry and the acceleration profile. For this, a custom hybrid global/local constrained optimiser based on Bayesian Optimisation and a gradient-less local algorithm has been developed with these kind of problems in mind.

This chapter is divided in three sections. Section 5.1 describes the catapult-assisted takeoff problem, together with its parametrisation. Next, section 5.2 provides an introduction to the Bayesian Optimisation theory, as well as a discussion of the implementation and the local refinement of the solution using the Gaussian Process model as a surrogate for it, while section 5.3 describes the cost function used to evaluate the performance of the manoeuvre. Numerical results can be found in section 5.4, which includes an assessment of elastic effects, an exploratory parametric study and the results from the optimization process.

## 5.1 Manoeuvre definition

To exemplify the previous approach we consider the optimization of a catapult-assisted takeoff on a representative VFA model. This section covers the description and parametrisation of the maneuver and the cost function definition.

### 5.1.1 Catapult-Assisted Take Off Parametrisation

The aircraft model used for this example is the T-tail HALE (see Section 3.6). It is trimmed for level flight at $10\,\mathrm{m\,s^{-1}}$ and the trim parameters for this condition are given in table 3.16.

HALE aircraft are rarely designed to sustain its own weight in the ground. Consequently, the aircraft is first attached to the catapult in multiple points of the wings and fuselage (these attachment points are also assumed to be rigidly linked among themselves, as they would be on a moving part of the catapult), then accelerated along a prescribed trajectory (a *ramp*), and finally released when a certain speed is reached. Only displacement constraints are enforced at the attachment points. A straight ramp is assumed, defined by its angle with respect to the ground, $\Gamma_R$. While some of the assumptions in this analysis may not apply to particular catapult designs, this problem is posed as a proof of concept. Its generalisation does not affect the validity of the methods here proposed. The aircraft is mounted on the catapult with an additional offset angle $\theta_o$, as shown in figure 5.1. Its attitude angle with respect to the horizontal is $\alpha_{\text{cruise}} + \theta_o$, with $\alpha_{\text{cruise}}$ the cruise angle of attack shown in table 3.16 (which will change depending on the flexibility of the model). A reference root angle of attack is defined while in the ramp by subtracting the ramp angle, $\Gamma_R$, to that attitude angle. A constant acceleration, $a$, is

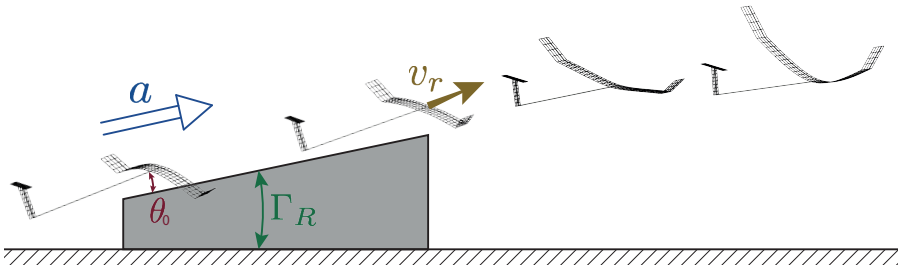



**Fig. 5.1** Parameters in the catapult-assisted takeoff. The aircraft goes from rest on the left to level flight on the right

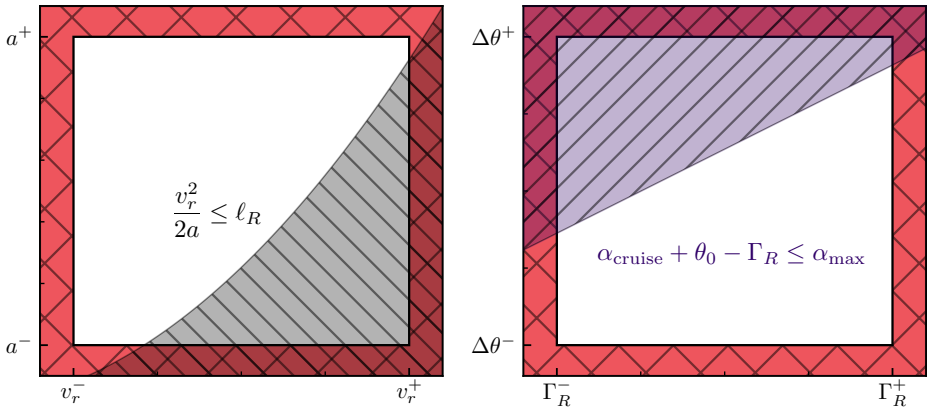



**Fig. 5.2** Parameter space with bounds in red and ramp length and angle of attack constraints in green and purple, respectively.

assumed through the takeoff. Lastly, the release speed of the aircraft is also a parametric value identified as $v_r$.

In addition to the previous four parameters and the assumptions that come implicitly with them (straight takeoff trajectory, constant acceleration), some constraints need to be imposed. Firstly, a maximum ramp length $\ell_R$ is defined by means of a unilateral (inequality) constraint in terms of acceleration and release velocity, as shown in figure 5.2. Secondly, the angle of attack of the aircraft during the launch is limited to be under the stall limit. This is enforced by a constraint on the incidence angle, $\alpha_{\text{cruise}} + \theta_o - \Gamma_R$ (Figure 5.2). The parametric domain that is studied in this work is therefore a combination of bounds for the four parameters plus the added constraints of the ramp length and incidence angle, which form a convex domain.

During takeoff, the aircraft is subject to the reaction forces at the clamping points. These reaction forces depend on the dynamics of the aircraft structure and the aerodynamic forces that change with the structural deformation and the increasing speed. When

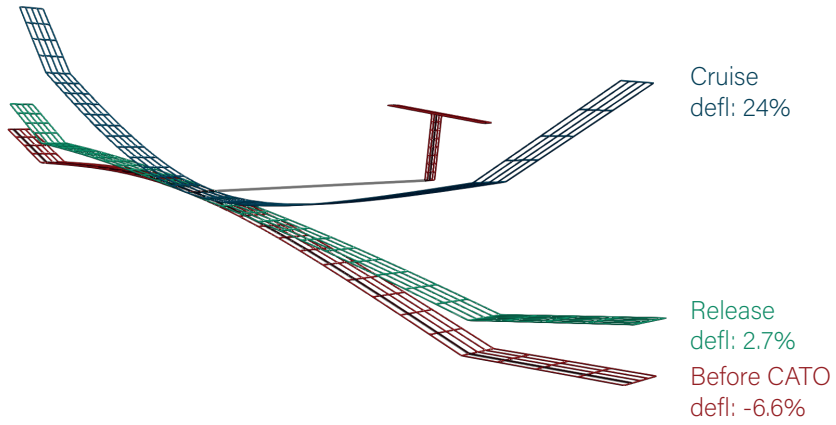



**Fig. 5.3** Snapshots of the aerodynamic grid and wing tip deflections at different stages of the baseline CATO.

| | |
|---|---|
| Acceleration, $a$ | $4.10\ \mathrm{m\,s^{-2}}$ |
| Release speed, $v_r$ | $10.77\ \mathrm{m\,s^{-1}}$ |
| Ramp angle, $\Gamma_R$ | $7.26°$ |
| Attitude offset, $\theta_o$ | $4.87°$ |

**Table 5.1** Baseline for parametric study

the aircraft is at rest, the reaction forces only counteract the gravity. As the aircraft is accelerated, the inertial and stiffness effects also need to be counteracted in order to enforce a smooth trajectory and acceleration profile. In the final stages of the CATO, the aerodynamic effects are dominant and the clamping forces act to keep the aircraft from deviating from the path by taking off too early, avoiding structural deformations between clamping points, and potentially fighting the increase in drag due to the lift generation. To exemplify this, figure 5.3 shows the instantaneous shape of the aircraft on a baseline CATO (see Table 5.1), together with the wing tip deflection (in percentage of the semispan) with respect to the jig shape.

The clamping point location is therefore an important factor in a catapult-assisted takeoff. As a rule of thumb, a catapult built with a shape which is close to that of the aircraft in the cruise state would result in smaller structural loads and variations in altitude. However, this would not be a practical solution, and our approach has been to study the effects of large structural deflection variations throughout the whole maneuver. This requires a more challenging numerical setup but a much smaller catapult. The results shown here are simulated with the aircraft supported in three points: on both sides of the wing root, separated 1 m from the root, and the fuselage-fin intersection. The release occurs simultaneously in all three points.

### 5.1.2 Simulation process

Simulation of the CATO is split in the three stages shown in Figure 5.4: *initial stage*, *launch stage* and *free flight*. During the initial and launch stages, the nodal velocity constraints are active, while the free-flight stage is a nonlinear aeroelastic simulation including rigid body dynamics. The initial stage defines the equilibrium shape of the aircraft on the catapult at rest under gravity loads. A background flow (incoming wind at ground level) velocity is assumed and $0.5\,\mathrm{m\,s^{-1}}$ has been set as a default. For convenience in our implementation, the initial equilibrium is found from a dynamic aeroelastic solution starting with undeformed airframe. During the solution, the aerodynamic wake is convected using the background flow velocity. As a result, the wake length during this *initial* stage is sensibly shorter than in free flight. This does not affect the solution due to the fact that the aerodynamic influence during this stage is negligible, so the shorter wake influence on the solution can be dismissed. However, the robustness of the vortex-induced velocity model, as described in section 2.2.5.1, is critical here because the distance between the wake vortex segments and the trailing edge can be small enough to cause numerical instability.

Once the structural state of the configuration has approached the steady-state deformation with the gravity and background flow loads, the launch stage begins. The three components of the velocity of the attachment points are prescribed based on the launch parameters as a function of time. This stage is the most demanding from a numerical robustness point of view. The first part of the launch is dominated by inertia effects, where the mass distributed along the high aspect-ratio wing is suddenly accelerated from a number of clamping points located near the wing root. In addition, as the wing generates more lift, its geometry changes, which can result in geometrical couplings that did not exist in the first stages of the launch. This change can be too quick for the aeroelastic modal frequencies relevant here (mostly OOP bending), and the state of the system can differ noticeably from the corresponding quasi-steady one. As the CATO progresses, dynamic pressure increases and aerodynamics play an increasingly important role, and the wake stretches to several chords downstream. Unsteady aerodynamics result in damped wing out-of-plane bending dynamics, while in-plane and torsion dynamics are geometrically coupled due to the effective dihedral and are relatively undamped. The release from the ramp is a critical point from a numerical stability point of view. The structural solver switches from a constantly accelerated motion enforced through nodal multibody constraints to an unsteady free flight simulation.

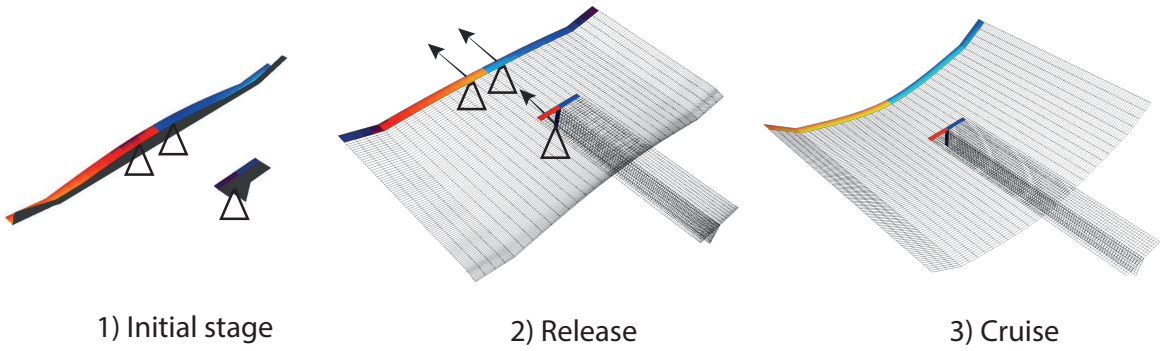



**Fig. 5.4** Three steps of the CATO process, showing an increasing wake length as the aircraft velocity increases.

After the release, the aircraft dynamics are simulated in open-loop, with the trim parameters set for cruise flight. This decision is due to two reasons: first, it reduces the complexity in the simulation and the analysis, and second, existing solar-powered HALE configurations generally do not feature high-authority flight control systems due to their very specific mission, payload and requirements. We note also that the typical climb rates are so small that optimizing launch for initial level flight or climb fundamentally gives the same outcome.

All simulations are run using a simplified convected wake for the UVLM where the wake is convected with the background flow and not by self-generated induced velocity (see Section 2.2.3). Finally, an investigation onto ground effects showed that they are negligible for such light vehicle and they have not been included in the simulations due to the high cost of including them in the UVLM formulation.

The mixed coupling formulation introduced in Section 2.3.3 increases the order of the approximation of the aerodynamic loads and provides more accurate results for a given time step when compared to a weakly coupled simulation. While the number of evaluations of the structural solver increases, a coarser aerodynamic spatial discretization and the superior accuracy of the strongly coupled simulations—specially when including rigid body motions that need to be integrated— result in substantial computational saving for highly dynamic cases. In particular, the model used in this work has second in-plane bending mode frequency of around 14 Hz in its cruise-deformed shape. Based on the chordwise discretization and the cruise speed, the FSI—and aerodynamic—time step for the simulation has been chosen $\Delta t = 2.5 \times 10^{-2}$ s, which equates to a Nyquist frequency of 20 Hz. While the second in-plane bending mode would be captured if the

structural solver was run with the same time step than the FSI, the accuracy of the beam dynamics would be compromised. With the in-plane bending component of the dynamics being of such importance for this particular problem, this simulation has been run with $\Delta t_{\text{struct}} = 0.5\Delta t_{\text{FSI}}$. This doubles the Nyquist frequency for the structural solver to 40 Hz, which has been deemed high enough to capture the relevant dynamics of the problem accurately.

### 5.1.3 Numerical considerations

#### 5.1.3.1 Imposition of nodal velocity constraints for CATO simulation

Enforcing the nodal trajectories for the CATO maneuver can be achieved in a number of ways, but the most practical ones involve prescribing the nodal velocities in $A$ or $G$ frame of reference at the clamping points. The non-holonomic constraint (see Section 2.1.3) for a prescribed nodal velocity $\boldsymbol{v}_{\text{node}}$ in $A$ frame of reference can be expressed as

$$\boldsymbol{\Phi}_{nh} = \dot{\boldsymbol{R}}_{\text{node}} - \boldsymbol{v}_{\text{node}} = 0 \tag{5.1}$$

$$\boldsymbol{B}_{nh} = \mathbb{I}_{3\times 3}. \tag{5.2}$$

While the non-holonomic constraint for a prescribed nodal velocity $\boldsymbol{v}^G_{\text{node}}$ in $G$ frame of reference is introduced as

$$\boldsymbol{\Phi}_{nh} = \boldsymbol{C}^{GA}\left(\boldsymbol{v}_{\text{FoR}} + \boldsymbol{\omega}_{\text{FoR}} \times \boldsymbol{R}_{\text{node}} + \dot{\boldsymbol{R}}_{\text{node}}\right) - \boldsymbol{v}_{\text{node}} = 0 \tag{5.3}$$

$$\boldsymbol{B}_{nh} = -\boldsymbol{C}^{GA}\text{skew}(\boldsymbol{R}_{\text{node}}) \tag{5.4}$$

$$\boldsymbol{K}^t = \frac{\partial \boldsymbol{C}^{AG}\dot{\boldsymbol{\lambda}}}{\partial \boldsymbol{p}} = -\boldsymbol{C}^{AG} + \text{skew}(\boldsymbol{R})\frac{\partial \boldsymbol{C}^{AG}\dot{\boldsymbol{\lambda}}}{\partial \boldsymbol{p}}. \tag{5.5}$$

With these two types of constraints, the CATO maneuver can be enforced by constraining as many nodes as clamping points exist and prescribing velocities in inertial frame of reference such that their relative positions remain unchanged (simulating a rigid rig). It can also be enforced in two stages: first, all the clamping nodes are constrained to have a velocity in $A$ frame of reference equal to 0, effectively creating a rigid link between them. Then, several options are available in order to prescribe the velocity of the $A$ frame of reference itself. For example, two nodes can have their $G$ velocity prescribed in order to control the attitude of the aircraft, or the $A$ frame of reference velocity and

rotation velocity can be prescribed independently. For this work, the first approach will be followed in order to simplify the setup process.

## 5.2 A Bayesian optimization framework for complex maneuver definition

If the number of independent parameters used to define the manoeuvre (i.e. its *dimensionality*) is sufficiently small, gradient-less methods become competitive in the exploration of the complex design space of feasible maneuvers, specially when the objective function is expensive to evaluate, and there are no numerical or analytical approximations available for the Jacobian or Hessian matrices. Many gradient-less algorithms, such as Powell [142] and Nelder-Mead [128] are still of local character and are designed to work sequentially, that is, one single cost function can be evaluated at a time, and the solution is dependent on the initial guess. Next, global gradient-less methods such as genetic algorithms require too many function evaluations and tuning its parameters is a long and problem-dependant process. A more suitable family of global optimisation algorithms is Bayesian Optimization (BO) [177, 55]. BO is now extensively used in machine learning for optimization of the hyperparameters of a learning model. It is best suited for expensive cost functions defined in continuous domains with a relatively small dimensionality (typically less than 20). BO is based on two steps per iteration: the construction of a surrogate for the cost function and the estimation of the uncertainty of that surrogate using Gaussian process regression. This information is then used to determine where to sample next using a balanced trade-off between exploration versus exploitation. This results in a global optimiser that samples the domain in a smart way while building a surrogate of the objective function on the fly.

Let $\boldsymbol{x} \in \mathcal{X} \in \mathbb{R}^n$ be a parametrisation of the space of possible maneuvers under consideration. Here, $\mathcal{X}$ is the specified domain and $n$ is the dimensionality of $\boldsymbol{x}$. An optimization problem is then defined with a cost function $J(\boldsymbol{x}) \in \mathbb{R}$. Adding a set of unilateral constraints $\boldsymbol{\Phi}(\boldsymbol{x})$, the problem can be written as

$$\min_{\substack{\boldsymbol{x}\in\mathcal{X};\\ \boldsymbol{\Phi}(\boldsymbol{x})\leq 0}} J(\boldsymbol{x}). \tag{5.6}$$

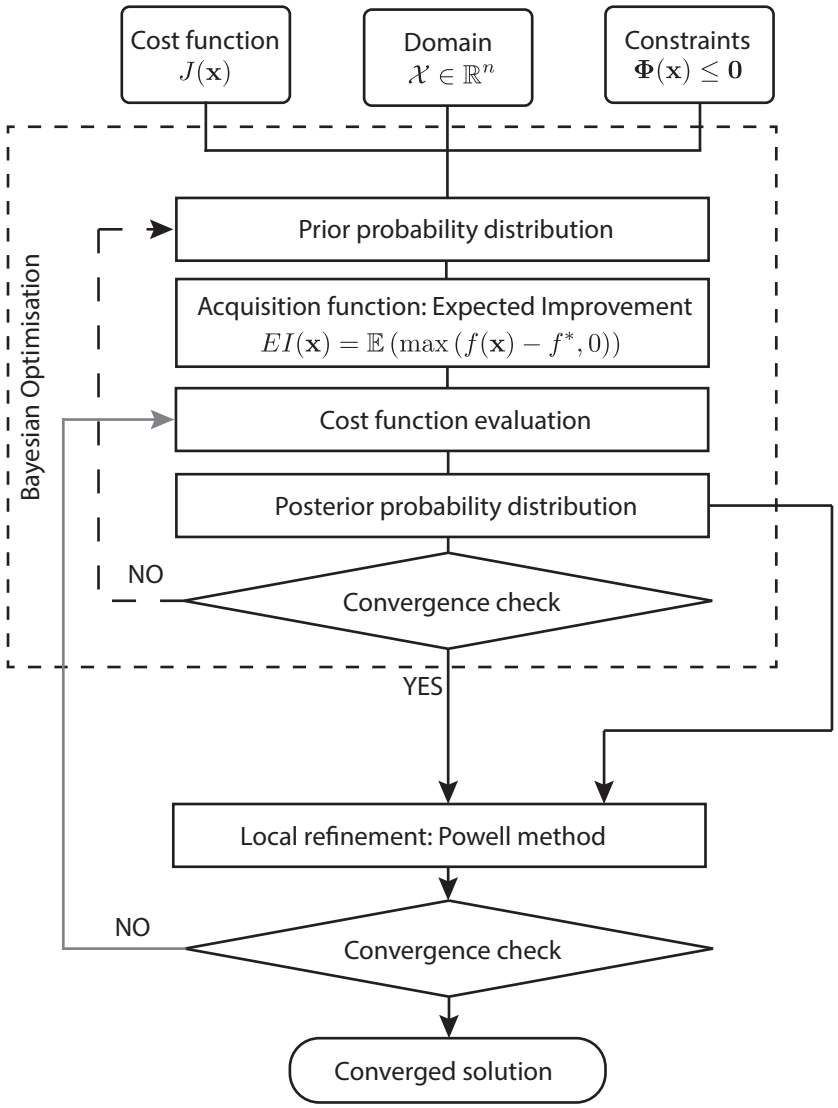



**Fig. 5.5** Flowchart of the hybrid global/local optimisation process.

Figure 5.5 shows a flowchart of the global optimisation algorithm based on Bayesian Optimisation, together with the local refinement step applied once the global algorithm convergence slows down.

### 5.2.1 Surrogate construction

As stated before, a BO framework requires two components: first, a *Bayesian statistical model* of the objective function, and second, an *acquisition function* that decides where the cost function will be sampled next [177]. Typically, the model is first initialized with a set of evaluations designed to cover the parametric space, such as a latin hypercube sampling.

However, since complex constraints are sought here, that structured initialization is not straightforward and the points are taken instead at random from inside the valid domain.

In broad terms, the whole Bayesian Optimisation algorithm can be described as shown in Algorithm 1.

**while** $n \leq N_{eval}$ **do**
  Update the posterior probability distribution on $J$ using all the available data;
  Let $x_n$ be a maximiser of the acquisition function over $x$, where the acquisition function is computed using the current posterior distribution;
  Observe $y_n = J(x_n)$;
  Increment $n$;
**end**
Return a solution: either the point evaluated with the largest $J(x)$, or the point with the smallest posterior mean;

**Algorithm 1:** Simple pseudo-code for BO, adapted from [55].

Gaussian Process (GP) regression is a statistical approach for function modelling, in this case, our objective function. Assuming we have a finite collection of points $x_1, \ldots, x_k \in \mathbb{R}^d$ and their respective $f$ values arranged in a vector $[f(x_1), \ldots, f(x_k)]$. In Bayesian statistics, a unknown quantity such as the vector of function values is supposed to be obtained from a random process with a probability distribution. The GP regression assumes the probability distribution to be a multivariate normal characterized by its mean vector and covariance matrix. The mean vector and covariance matrix are built evaluating the *mean function* $\mu_0$ at every point $\boldsymbol{x}_i$ and the *kernel* $\Sigma_0$ at every pair of points $(\boldsymbol{x}_i, \boldsymbol{x}_j)$, respectively. The kernel is a function that reflects the belief that points that are closer together have a similar function value than others separated farther apart.

Two important concepts in Bayesian statistics are the *prior* and *posterior* probability distributions. The first one is the probability of an event *before* new observations have been made, that is, it is our best guess about a probability based on our current knowledge. On the other hand, a *posterior* probability distribution is the revised probability distribution *after* considering new information (in our case, evaluating the objective function). A posterior probability distribution is computed by updating the prior distribution using Bayes' theorem [147]:

$$P(A|B) = \frac{P(A \cap B)}{P(B)} = \frac{P(A) \cdot P(B|A)}{P(B)}. \tag{5.7}$$

For our problem, the *prior* probability distribution given a vector of values of $f$ for $n$ samples: $f(\boldsymbol{x}_{1:n}) = [f(\boldsymbol{x}_1), \ldots, f(\boldsymbol{x}_n)]$ is

$$f(\boldsymbol{x}_{1:n}) \sim \mathcal{N}\left(\mu_0(\boldsymbol{x}_{1:n}), \Sigma_0(\boldsymbol{x}_{1:n}, \boldsymbol{x}_{1:n})\right). \tag{5.8}$$

With $\boldsymbol{x}_{1:n}$ observed for a certain $n$, we want to infer $f(\boldsymbol{x})$ for a new point $\boldsymbol{x}$. This next evaluation is noted $k = n + 1$. Then, the prior over $[f(\boldsymbol{x}_{1:n}), \boldsymbol{x}_k]$ is calculated with equation (5.8). Using Bayes' rule, the conditional distribution of $f(\boldsymbol{x}_k)$ given the previous observations (that is, the *posterior probability distribution*) is

$$f(\boldsymbol{x})|f(\boldsymbol{x}_{1:m}) \sim \mathcal{N}\left(\mu_n(\boldsymbol{x}), \sigma_n^2(\boldsymbol{x})\right) \tag{5.9}$$

$$\mu_n(\boldsymbol{x}) = \Sigma_0\left(\boldsymbol{x}, \boldsymbol{x}_{1:n}\right) \Sigma_0\left(\boldsymbol{x}_{1:n}, \boldsymbol{x}_{1:n}\right)^{-1} \left[f(\boldsymbol{x}_{1:n}) - \mu_0(\boldsymbol{x}_{1:n})\right] + \mu_0(\boldsymbol{x}) \tag{5.10}$$

$$\sigma_n^2(\boldsymbol{x}) = \Sigma_0\left(\boldsymbol{x}, \boldsymbol{x}\right) - \Sigma_0\left(\boldsymbol{x}, \boldsymbol{x}_{1:n}\right) \Sigma_0\left(\boldsymbol{x}_{1:n}, \boldsymbol{x}_{1:n}\right)^{-1} \Sigma_0\left(\boldsymbol{x}_{1:n}, \boldsymbol{x}\right). \tag{5.11}$$

With a statistical model of the objective function, only the acquisition function needs to be defined. In this work, the Expected Improvement (EI) [193] criterium is used. It is defined as

$$EI(\mathbf{x}) = \mathbb{E}\left(\max\left(f(\boldsymbol{x}) - f^*, 0\right)\right), \tag{5.12}$$

where $\mathbb{E}$ is the expected value given the posterior probability distribution, and $f^*$ is the maximum value of $f$ evaluated so far.

Once the BO has approached the solution and the function has been evaluated around the global minimum, the uncertainty of the Gaussian-process surrogate in that region is very low and therefore other points of the parametric space are sampled instead. This, which is one of the strengths of BO, turns out to also slow down convergence in the final stages of the optimization process. An alternative strategy for local refinement has been developed, and is described next.

### 5.2.2 Local refinement of the solution

One common algorithm used in global optimisation is called Basin-Hopping [198]. The original method is based on random perturbation of the global coordinates to initialize the local optimizer which then provides a value for the *local*–potentially global– minimum. While this method is suitable for rugged "egg-tray-like" but relatively cheap objective functions, it is not readily applicable for our smoother and more expensive case. Building on the idea of combining global and local optimisation techniques, the strengths of the

GP regression technique have been leveraged and the weakness have been minimized. In the *global* stage, the BO framework has been used to smartly sample the domain and approximate the local minimum.

Once the *global* optimization described above has found a potential minimum, a *local* method is used to refine the solution [198, 83]. Typical basin-hopping approaches require either a rather large number of additional function calls or Jacobian/Hessian information. As our cost function is expensive, we use instead the information contained in the Gaussian process regression of the BO as a surrogate of the cost function. Moreover, the BO framework is used as an adaptive sampling algorithm to increase the accuracy of the surrogate near the optimum.

After building this surrogate from the trained Gaussian process, we need to add the constraints so that the local optimizer does not converge towards a point outside the domain. Local optimisers that support constraints are usually more expensive or present slower convergence rates. Instead, the bounds have been externally enforced as a penalty. The local optimizer used in this work is the Powell's Conjugated Direction Method [142] implementation in SciPy, a local gradient-less optimisation algorithm. While gradient-based algorithms are available, and usually outperform non-gradient based ones, enforcing constraints through a penalty approach generates an artificially abrupt step in the value of the function in our problem. This then results in convergence problems as the Jacobian is estimated through finite difference methods. It is worth pointing out that if no constraints had to be introduced–or the minimum was relatively distant from them–, a gradient-based method can be used with the Jacobian estimation natively provided by the Gaussian process surrogate.

### 5.2.3 Implementation

The implementation of the BO optimizer used is the one distributed in GPyOpt [186], a separate framework that deals with Gaussian Process Regression Optimisation. This framework supports simple bounds constraints in the parameters, as well as unilateral constraints expressed as inequalities based on the input parameters. In addition, the objective function evaluations can be performed in parallel with just a small penalty in the convergence rate of the algorithm. The call to SHARPy is done in a manner that avoids race conditions that could potentially arise from the parallel execution of the cost function evaluation. The input/output for every case is handled independently, from separate folders, by parallel subprocesses running an instance of SHARPy each.

GPyOpt data structures allow direct access to the GP regression model, which is then used in the local improvement of the solution by means of a wrapper that includes the constraint terms. This wrapper is then fed to the SciPy Powell optimiser as the objective function. A diagram of the whole process is shown in Figure 5.5.

## 5.3 Optimal maneuver definition

Once the problem has been defined, it is necessary to provide a way of quantifying the suitability of the different maneuvers based on the values of the parameters.

While multiple cost functions can be defined for this maneuver and several were investigated in preparation of this work, *the selected cost function will minimise a combination of ground clearance after takeoff and structural root loads*, for different (prescribed) available runway lengths.

The ground clearance is quantified as the minimum value of wing root vertical coordinate with respect to the ground $z = 0$ plane evaluated after release. While low —or negative— values of ground clearance can obviously be considered bad performance, relatively high values of ground clearance past a certain point do not improve the suitability of the maneuver anymore. Due to this, the function chosen to map the ground clearance to a penalty term in the cost of the takeoff is a sigmoid such as

$$J_{\text{ground}} = \frac{S_{\text{ground}}}{1 + \exp\left(\tilde{z} - \frac{1}{2}\right)}, \tag{5.13}$$

with $\tilde{z} = (z - z_{\min})/(z_{\max} - z_{\min})$, and $z_{\min}$ and $z_{\max}$ given constants with the minimum and maximum clearances described previously (their effect on the cost function can be seen in Figure 5.6). Finally, $S_{\text{ground}}$ is a scaling factor in the cost function to balance both contributions.

The contribution due to structural loads to the overall cost is normalized based on the $10\,\mathrm{m\,s^{-1}}$ cruise loads for the same aircraft. Three wing root loads are considered: torsion, out-of-plane bending and in-plane bending moment. Being $\mathbf{M}_{\text{ref}}$ the moments of the wing root in cruise at $10\,\mathrm{m\,s^{-1}}$, and $\mathbf{M}$ the instantaneous root loads at a time $t$, the contribution of the wing root loading to the cost function is defined as

$$J_{\text{loads}} = \sum_{k=1}^{3} \max_{t} \left( \left| \frac{M^k}{M^k_{\text{ref}}} \right| - 1 \right) \cdot S^k_{\text{loads}} \tag{5.14}$$

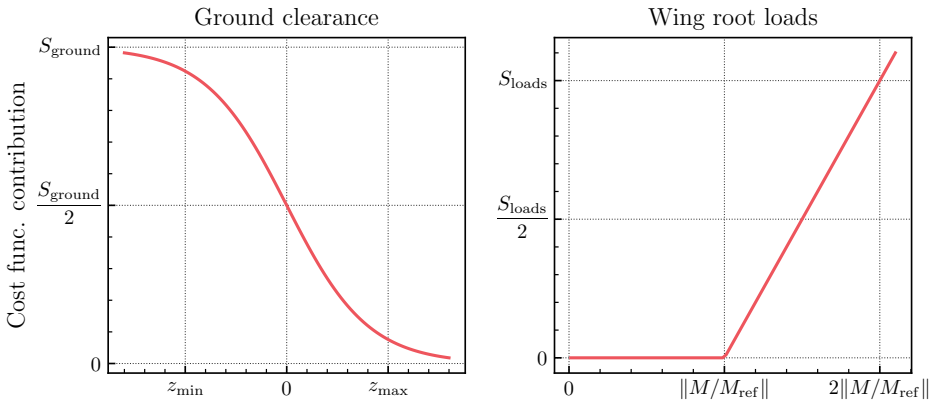



**Fig. 5.6** Cost function contributions associated to the ground clearance (left) and wing root loads (right).

where $\boldsymbol{S}_{\text{loads}} \in \mathbb{R}^3$ is a scaling vector. The overall shape of this contribution is shown in figure 5.6. The final cost function includes both contributions, that is, $J = J_{\text{ground}} + J_{\text{loads}}$. Despite the hybrid optimiser user here, this whole process is a standard multiobjective optimisation process where ground clearance and wing root loads are the measures of performance, the launch parameters (release velocity, acceleration, ramp angle and attitude offset angle) are the design variables, and the ramp length and maximum effective incidence angle are unilateral constraints imposed on the launch parameters.

## 5.4 Numerical Results

### 5.4.1 Flexibility effects on a catapult-assisted takeoff

Baseline CATO parameters have been established after a preliminary numerical exploration and they are given in table 5.1. It corresponds to a runway length $\ell_R = 14.15\,\text{m}$ and incidence angle $\alpha_{\text{cruise}} + \theta_0 - \Gamma_R = 1.68°$. For this maneuver, the two T-tail variants presented in Section 3.6 are considered: the *flexible* (baseline), and the *stiff*, both defined in table 3.16. Figure 5.7 shows the longitudinal trajectory of the wing root node as seen from an Earth system. From a ground clearance point of view, the very flexible aircraft presents good performance, as its lowest point after release is effectively the release point. However, the differences with the stiff aircraft are stark. With the given maneuver definition, the stiff aircraft is released near its steady-state condition, with just a small pitch angle deviation from the cruise conditions and a slightly greater speed than the original cruise speed. Shortly after release, the response of the stiff aircraft is

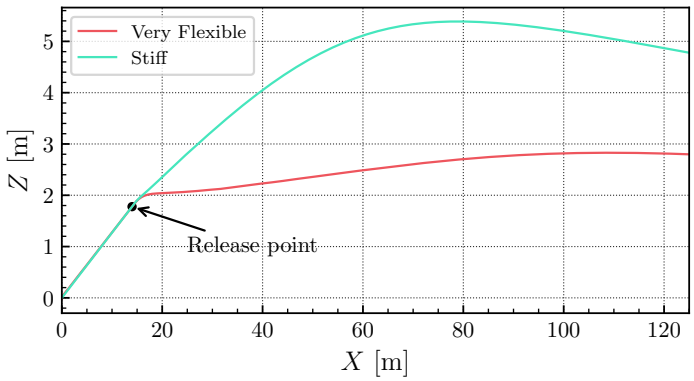



**Fig. 5.7** Longitudinal trajectory during CATO of the T-tail HALE in its very flexible and stiff variants. The parameters of the launch can be found in table 5.1.

dominated by its phugoid mode, which is slow enough that it could be further damped using a very basic flight controller.

Figure 5.8 presents the three components of the wing root moments, normalized by the $10\,\mathrm{m\,s^{-1}}$ cruise root loads of the very flexible aircraft (for the flexible and stiff aircraft results). The first 1.5 s in the simulation correspond to the initial stage in figure 5.4, where the aircraft finds its equilibrium on a stationary catapult. The out-of-plane (OOP) bending moment therefore drops for both configurations. During the launch stage this moment can be seen to increase much faster on the stiff aircraft than on the very flexible one. This is mainly due to two facts: first, the stiff configuration gravity-only shape is very similar to the jig and cruise shapes, and second, the stiff aircraft has quicker dynamics (higher natural frequencies). It is also important to note the large damping of the out-of-plane oscillations due to the aerodynamics, together with its very limited coupling with the other degrees of freedom. The torsion and in-plane (IP) dynamics oscillate at the same frequency with no phase delay due to geometrical coupling. In addition, the oscillatory behavior of these two degrees of freedom is only lightly damped. This is due to the much smaller aerodynamic damping in these motions when compared to the out-of-plane bending. It is important to note that, unlike other common aerodynamic solvers such as DLM, the aerodynamic solver used in these results accounts for the incident velocity contribution of the in-plane motion of the structure. This includes lift as well as induced drag variations. The details on the process followed for obtaining the aerodynamic forces from the UVLM solutions in a more general way than the usual method described in [89] can be found in [171]. Moreover, although the steady flight deformations of the models are substantially different, the wing root bending loads are

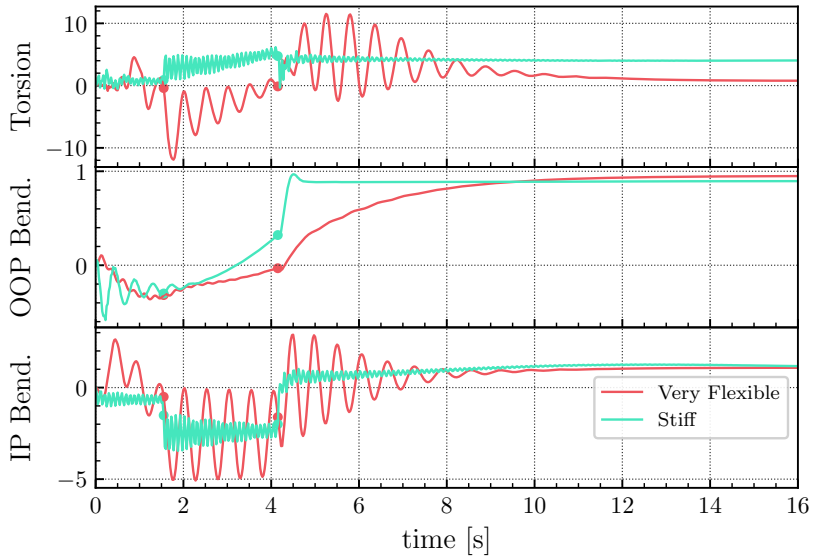



**Fig. 5.8** Torsion, out-of-plane and in-plane wing root loads during CATO for the very flexible and the stiff T-tail models. The parameters of the launch can be found in table 5.1. Dots indicate start/end of launch stage and values are normalized with cruise wing root loads.

not dissimilar, as they are driven by the total weight distribution, which is the same in both models. The differences in torsion loading are more marked due to the large effective dihedral of the very flexible aircraft, which results in a different compensation of the induced drag component of the aerodynamic forces. In fact, the wing tip structural element of the stiff configuration has a dihedral angle of 20°, the jig shape one, while the very flexible configuration presents an effective dihedral angle of 38°.

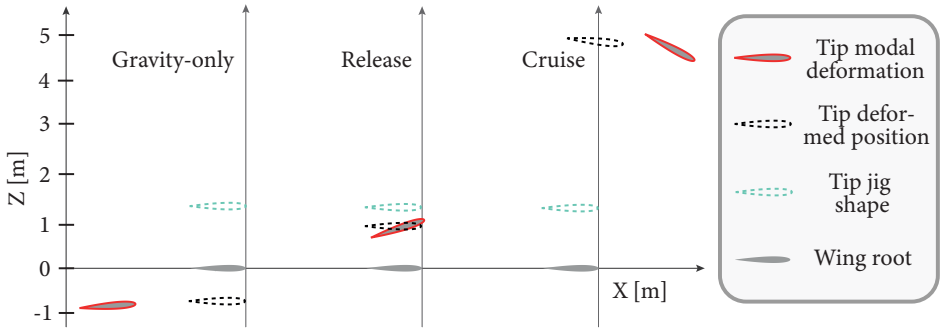



**Fig. 5.9** Position of the wing tip cross section of the flexible model at different stages of the CATO. Position for the jig shape, aeroelastic shape and second vibration mode with the former as reference condition.

Lastly, in order to better understand the variation of the characteristics of the structure during the different stages of the takeoff, the first natural vibration modes

| Mode | Gravity-only | | Release | | Cruise | |
|---|---|---|---|---|---|---|
| | **Freq [Hz]** | **Type** | **Freq [Hz]** | **Type** | **Freq [Hz]** | **Type** |
| 1st mode | 0.449 | OOP (I) | 0.438 | OOP (I) | 0.419 | OOP (I) |
| 2nd mode | 1.83 | IP/torsion | 1.81 | Torsion/IP | 1.34 | IP/Torsion |
| 3rd mode | 2.72 | OOP (II) | 2.71 | OOP(II) | 2.61 | OOP (II) |

**Table 5.2** First three natural vibration frequencies of the very flexible model at different stages of the CATO.

(without aerodynamics) of the very flexible aircraft have been calculated from the tangent mass and stiffness corresponding to three different conditions, namely, the aircraft at rest (gravity only), at the release point, and in cruise. The results at the release point have been calculated with the non-equilibrium conditions as reference. Table 5.2 shows the frequencies of the first three vibration modes, together with their characterization. While out-of-plane bending modes do not fundamentally change with the configuration, the second mode in the table, which is a coupled in-plane/torsion mode, shows large changes. This is further investigated in Figure 5.9, which shows the projection on the plane of symmetry of the aircraft of the wing tip section for this second vibration mode. A normalized mode shape is superimposed on the tip deformed position, which is also included in the figure. As it can be seen, the second natural vibration mode is dominated by the in-plane bending component at rest, becomes mostly a torsional mode at the release point, and it clearly displays both contributions in cruise.

### 5.4.2 A parametric study

Prior to the actual optimization, the 4-dimensional parametric domain that describes the CATO of the very flexible aircraft is first explored by modifying each parameter with respect to the baseline defined in Table 5.1. Figures 5.10 and 5.11 show, respectively, the longitudinal trajectory and the wing root loads cost contribution calculated with the scaling parameters $\boldsymbol{S}_{\text{loads}} = [0.2, 0.2, 0.2]^\top$, when each parameter is independently modified from its reference value. The four parameters used in this analysis (acceleration, release speed, attitude offset and ramp angle) are divided in two groups: First, the parameters concerning the acceleration law—that is, acceleration and release speed—, and second, the parameters affecting the spatial trajectory and orientation of the aircraft during the CATO, namely, the attitude offset and ramp angle.

It can be observed how some parameters affect the scattering of the solutions in figure 5.10 much more than others. For example, trajectories obtained by changing the ramp

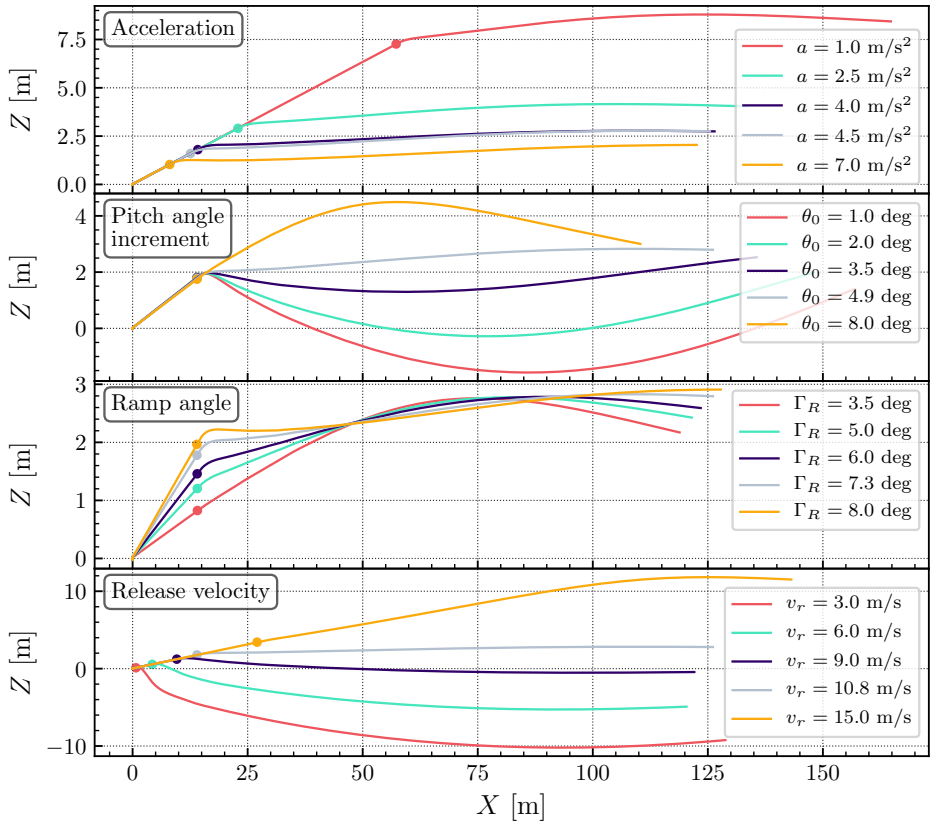



**Fig. 5.10** Wing root trajectories of the very flexible aircraft for changes to takeoff parameters with respect to the baseline (table 5.1). Dots indicate start/end of launch stage.

angle are close to overlapping, while the solutions with changing attitude offset, $\theta_o$, are essentially different. This might seem counter-intuitive, as they both control the attitude angle during the CATO, but the difference is that $\theta_o$ has direct control over the incidence angle of the wing during the whole manoeuvre. The main effect of flexibility during takeoff is that the wing out-of-plane dynamics is relatively slow compared to typical launch characteristic times. In this case, depending on ramp length constraints (we consider from 5 to 20 m), the typical frequency of the CATO is between $\omega_{\text{CATO}} \in [2.5, 4.4]$ rad s$^{-1}$ —longer ramps typically result in longer characteristic launch times—, while the first bending mode calculated with the release structural deformation as reference conditions is $\omega_{\text{OOP (I)}} = 2.75\,\text{rad}\,\text{s}^{-1}$. In the cases shown here, the OOP bending dynamics are in general slower than the characteristic time of the CATO launch for all but the longer ramp lengths, in which case, the time scales are similar. If the aerodynamic loads on the surfaces do not act early enough, or they are too small, the wings might not reach a state close to the cruise flight at release time. The acceleration and release speed

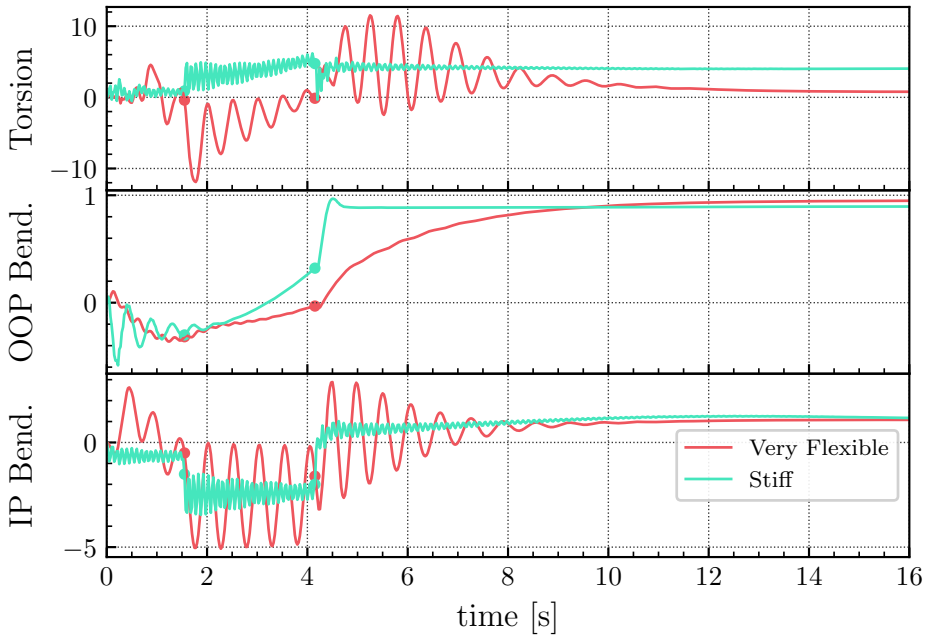



**Fig. 5.11** Instantaneous loads cost, $J_{loads}(t)$, for changes to individual takeoff parameters from baseline. Dots indicate start/end of launch stage.

parametric study shows that from a ground clearance point of view, the acceleration only affects the trajectory by a vertical offset, as the longer an angled ramp is (i.e., lower acceleration for the same release speed), the higher the release point is. However, the release speed affects the trajectory substantially, as shown in figure 5.10. This is due to two factors. The most direct one is that the aircraft is trimmed for cruise flight at $10\,\mathrm{m\,s^{-1}}$, and if it is released at a lower speed, it will lose altitude before reaching cruise speed, and thus, the ground clearance will be reduced. Second, higher release velocities result in higher dynamic pressure during takeoff, which deflects the wing to a close-to-cruise state faster. If at the release instant the structural deformation is not close to the cruise shape, the centre of mass of the aircraft, which is near the fuselage-wing intersection, will move downwards until the wings are generating enough lift and thus their shape is close to the steady flight one. Figure 5.12 shows the deformed shapes at release for release velocities.

Figure 5.13 shows the values of the cost function for the four-parameter study. First, a visible difference in variability of the cost can be observed. While the ramp angle $\Gamma_R$ has a reduced effect in ground clearance and minimal effect in loads, the release speed affects both contributions by a larger magnitude than any of the other parameters for the range of parameters here considered. What this figure does not capture is the possibly

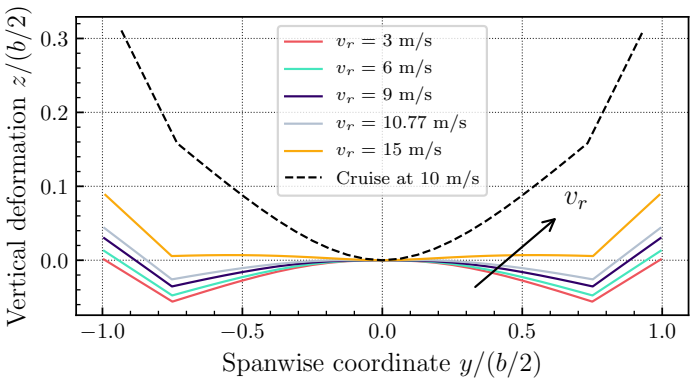



**Fig. 5.12** Wing shape at the release instant under changes of release speed from baseline.

more complex interaction of the attitude offset angle with the ramp angle with respect to the loads contribution of the cost. A direct study such as this would involve a large number of function evaluations, and this will be left to the optimization in the next section, which will capture, and exploit, this potentially complex behavior.

The interplay between acceleration and release speed is of special interest here. Figure 5.13 shows that greater release speed corresponds to a lower loads contribution to the cost function, but this is because, as the acceleration is kept fixed, the ramp length also increases. Once a ramp length constraint is introduced, the release speed is linked to the acceleration. This results in more abrupt accelerations for the high release speed cases, which increase the load contribution and also potentially increase the ground clearance contribution, as the slow out-of-plane wing dynamics might no have enough time to accommodate to the rapidly increasing dynamic pressure. Factoring this new coupling in the results is more complex than for the interaction between the attitude offset and ramp angles due to the acceleration effect on the loads. In addition, as the dynamic pressure increases—more or less rapidly depending on the acceleration and release speed parameters—, the influence of the attitude offset increases, as the aerodynamic loads will be greater, just as the damping introduced through added mass effects.

### 5.4.3 Optimization results

After the previous exploration, the maneuver parameter bounds have been chosen as in table 5.3. These values encompass a wide enough range so that the solution with a long ramp constraint is not bound by these limits. Optimal maneuvers are obtained for three fixed ramp lengths of 5, 10 and 20 m. The stall constraint in the incidence angle

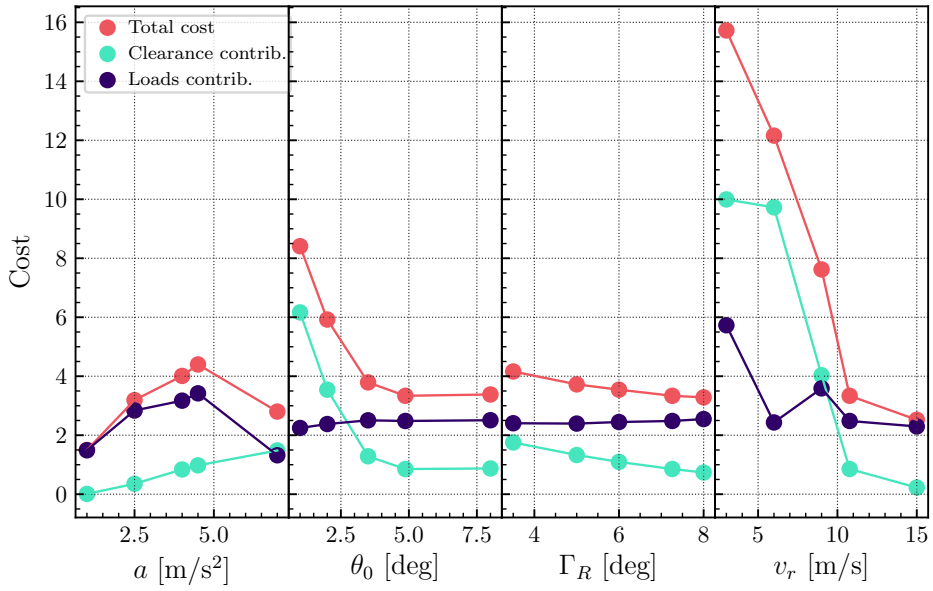



**Fig. 5.13** Contributions and total cost for a parametric study of the CATO maneuver.

| **Variable** | **Domain** |
|---|---|
| Acceleration, $a$ [m s$^{-2}$] | [1.0, 7.0] |
| Release speed, $v_r$ [m s$^{-1}$] | [3.0, 15.0] |
| Ramp angle, $\Gamma_r$ [°] | [-1.0, 8.0] |
| Attitude offset, $\theta_0$ [°] | [-1.0, 8.0] |

**Table 5.3** Bounds for parameters of the CATO optimization problem.

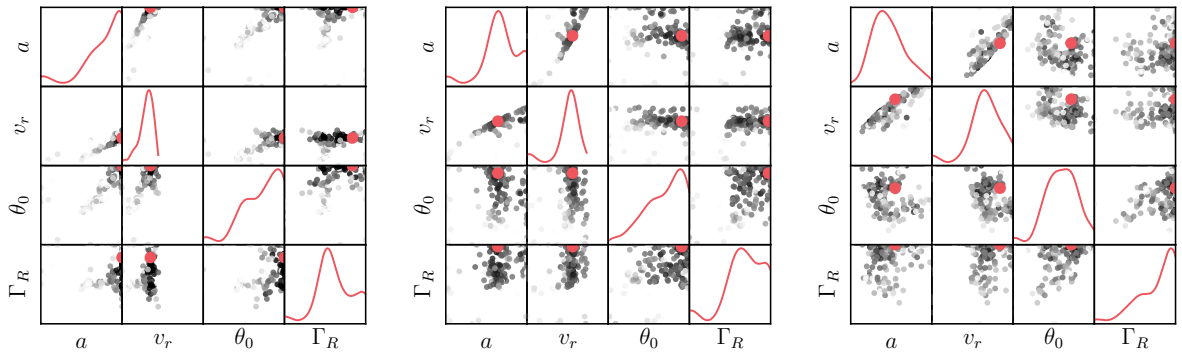
was active during the optimization. However, since it was not a relevant factor in the preliminary studies, it has been left fixed at 9°.

The cost function in this work, given in equations (5.13) and (5.14), has a number of parameters that need to be selected. First, $z_{\min}$ and $z_{\max}$ have been assigned -4 and 2 m, respectively. These values ensure that the slope of the cost function is not too small in the $z < 0$ domain so that convergence is accelerated. The $z_{\max}$ parameter has been chosen so that past 2 m there is not an important difference in cost function value to avoid the optimizer converging to solutions with high ground clearance but high wing root loads. Interestingly, if this saturation of the ground clearance cost is not implemented, the incidence angle constraint can be necessary to keep the simulation realistic. The ground clearance cost scaling parameter has been fixed at $S_{\text{ground}} = 10$ in order to obtain a cost of order one. Second, the wing root loads contribution $\boldsymbol{S}_{\text{loads}}$ has been assumed to be of the shape $\boldsymbol{S}_{\text{loads}} = s_{\text{loads}}[1, 1, 1]^T$ with $s_{\text{loads}} \in \mathbb{R}$ so that a single parameter needs to be

tuned. A value of $s_{\text{loads}} = 0.2$ has been chosen to balance the ground clearance with the wing root loads contributions. This means that both cost contributions are comparable in magnitude and thus a single contribution will not dominate the optimization.

Figure 5.14 shows a scatter matrix representation of the domain sampling for the three ramp length constraints in this work. The 4-dimensional parametric domain is represented as a $4 \times 4$ array of scatter plots. Each individual plot summarizes the relation between two variables while omitting the rest. The domain sampling shown in figure 5.14 corresponds to the Bayesian Optimization framework being run for 30 iterations with 5 parallel function calls per iteration, that is, 150 samples per figure. Dots in the scatter plot have a variable intensity that gives an indication of the associated cost of that sample. Here, darker spots represent better solutions with lower associated cost. The diagonal plots are a representation of the sampling density in the domain of the variable in the horizontal axis. From this representation, some insight into the cost distribution and the constraints effect can be drawn. First, the ramp length constraint, expressed as a function of $a$ and $v_r$ can be identified in the $a$ versus $v_r$ plots as a defined diagonal line that divides the domain in sampled and non-sampled. The ramp length constraint has a strong effect in the overall distribution of the samples and the optimum location in the domain. It is easy to see how the shorter the maximum ramp length is, the more the optimum tends towards a corner of the domain. This is true for all the variables, however the ramp angle $\Gamma_R$ behaves differently, with the darker sample dots, representing lower associated cost solutions, distributed quite horizontally across its domain. The fact that the ramp angle does not have an important effect in the CATO performance was already implied in section 5.4.2. The incidence angle constraint can be seen in the $\Gamma_R$ versus $\theta_o$ plot of figure 5.14a. The visible trend concerning this constraint is that the optimum approaches it when the ramp length constraint is very restrictive. From a physical point of view, it means that the aircraft increases the wing loading through incidence angle, with a low ramp angle and a high angle of attack, such that the wing deforms and approaches the cruise deformation as quickly as possible. When the ramp length is long enough, this constraint is not necessary, as a maneuver with such high incidence angle causes very large structural loads.

Figure 5.15 shows the longitudinal trajectories of the aircraft wing root node in the optimal takeoff corresponding to each each ramp length, $\ell_R$. Note that for $\ell = 5$ m the minimum altitude in the trajectory is negative, which would require a minimum height of the catapult to avoid having the aircraft touching the ground. Table 5.4 shows the CATO parameters of these optimum cases, together with their cost. First, the clearest trend




**(a)** Ramp length = 5 m. **(b)** Ramp length = 10 m. **(c)** Ramp length = 20 m.

**Fig. 5.14** Scatter matrix of the BO parameter sampling for three ramp lengths. Diagonals show sampling density, darker dots show lower cost, and red dot shows optimum.

| $\boldsymbol{\ell_R}$ [m] | $\boldsymbol{a}$ [m s$^{-2}$] | $\boldsymbol{v_r}$ [m s$^{-1}$] | $\boldsymbol{\theta_o}$ [°] | $\boldsymbol{\Gamma_R}$ [°] | **Cost** |
|---|---|---|---|---|---|
| 5 | 7.00 | 7.20 | 8.00 | 6.55 | 9.19 |
| 10 | 4.84 | 9.73 | 7.14 | 7.77 | 2.66 |
| 20 | 4.28 | 13.07 | 5.43 | 8.00 | 1.72 |

**Table 5.4** Results of the CATO procedure optimization for different ramp length constraints and loads-associated cost scaling.

is the greater ground clearance with longer ramps. As pointed previously, this is due to the longer launch times, which for a given release speed result in slower acceleration parameters. The very slow out-of-plane dynamics of the wing structure, around 0.45 Hz (see table 5.2) require long launch times. The constraint in ramp length, involving acceleration and release speed parameters introduces a dependency with consequences not shown in the previous parametric study, where the acceleration could be as low as possible with little change to the trajectory with respect to the release point (figure 5.10). The difference between those results and the ones in figure 5.15 is the ramp length limitation (see figure 5.2).

One important point worth mentioning is the influence of the shape of the ground clearance cost contribution, equation (5.13). Given that the function tends to zero when $z > z_{\max}$, the influence of the ground clearance in the overall result depends on the lowest altitude that the plane reaches from release until the end of the simulation. The $\ell_R = 20$ m case is released with $z > z_{\max}$, and it never dips under that value. That essentially means that the optimization for this problem is basically a loads minimization—as long as the plane never dips under $z_{\max}$, a case that would be instantly discarded by the optimizer—independently of the longitudinal trajectory that the aircraft follows. When

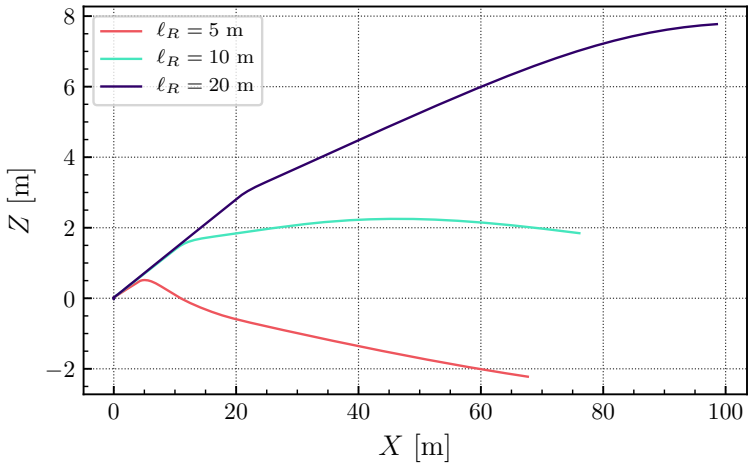



**Fig. 5.15** Longitudinal trajectories for the optimal solutions (Table 5.4) given different ramp length constraint values imposed as unilateral constraints. It can be seen how the $\ell_R$ =5 m case cannot achieve a positive ground clearance. This is an indicator than a longer ramp is a necessary requirement for this platform.

the constraint is more demanding, for example, $\ell_R = 5$ m, the ground clearance is the driving criterion for the optimizer, as the associated cost grows rapidly with low $z$ values. The medium case, $\ell_R = 10$ m is an interesting one, as the release point is close enough to $z_{\max}$ to be affected with the flattening of the cost function, while at the same time, any dip under this value will make the ground clearance contribution grow quickly. In figure 5.15 it can be seen how the optimum launch for this case is one that never goes too far above or below $z_{\max}$. The reason is that, if the trajectory went higher than this point, there would be no benefit in terms of ground clearance cost, while the loads would most likely increase. Likewise, if any point of the trajectory falls far below $z_{\max}$, the ground clearance cost will start to increase, decreasing the performance of the candidate.

From a loads point of view, the optimum launch procedures show a number of common features, which can be seen in the root moments shown in figure 5.16. First, the normalized loads in OOP bending are much lower than the torsion and IP bending ones. Given the mostly forwards acceleration suffered by the aircraft during the launch and the very low frequency of the first OOP bending mode, this is predictable. The high-frequency content in the OOP signal is linked to the second bending mode, which is excited in the initial transient. While conventional (linear) aeroelastic analysis does not focus on IP dynamics, the characteristics of this maneuver, together with the geometric coupling between torsion and IP that varies as the structure deforms during the launch, make the in-plane motion of the wing an important part of the analysis. Second, from

this geometric coupling, we can see how the torsion and IP loads present similar frequency responses during the launch phase. The launch phase goes from $t = 1.75$ s to $t = 2.6$, 3.9 and 4.7 seconds respectively for increasing ramp length, which can be easily seen on the time trace of all three root moments. Aerodynamic damping however is very low on the in-plane bending dynamics of the wing and has a much more dominant effect on the torsional response. This is due to the lack of first order aerodynamic effects acting on IP deformations, as they only generate an increase in incidence velocity. The torsion loads instead are aerodynamically damped due to the effect on the rotational velocity of the wing cross-sections, and their influence in the incidence angle of the wing has a coupling effect with the very damped OOP dynamics. It is also important to note how quickly the loads converge to the cruise values. This is due to the effect the long fuselage has on the damping of the longitudinal rigid body dynamics of the aircraft.

Amongst the differences between different ramp length cases, it is worth mentioning the different values of IP bending peaks during the launch phase. While the difference between the 10 and 20 m cases is not too large, the 5 m case has a peak IP bending loads value of up to 8 times the cruise IP bending loads. Despite this, up to release, the IP oscillations are of the same frequency, which is the 2$^{\text{nd}}$ structural mode indicated in table 5.2 and depicted in figure 5.9. However, past the release point, the motion becomes more damped and the frequency increases, which is consistent with the change of frequency that the second structural mode suffers when the wings of the aircraft start deforming due to the aerodynamic loading. Concerning the torsional dynamics, it is interesting to note how the peak value for all the procedures, but specially for the shorter one, is more than 10 times the cruise loads. The release instant can be identified in the torsion loads due to the higher frequency of the response, it shows an abrupt change of slope followed by a high-frequency oscillation. This is the kind of high-frequency-driven dynamics that needs a finer time discretization in the structural solver.

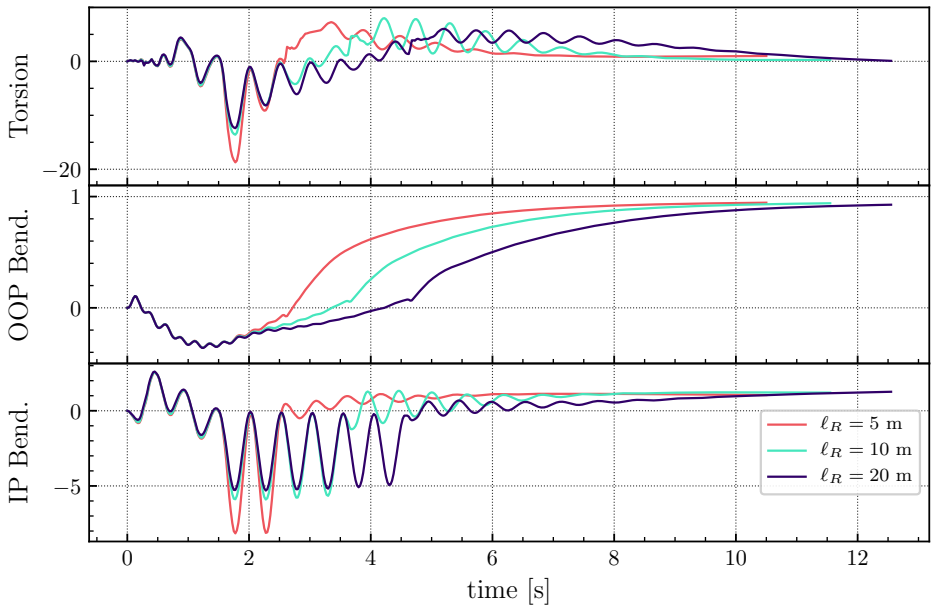



**Fig. 5.16** Wing root loads for the optimum cases (Table 5.4) given different ramp lengths. Loads normalized with cruise wing root loads.

# Chapter 6

# Conclusions

This chapter brings to an end the dissertation. Section 6.2 provides a summary of the most important outcomes of this work. Following this, section 6.1 describes the main contributions of this work, both in the theoretical and practical realms. Finally, section 6.3 identifies pending questions and proposes new research questions that have arisen from the outcomes of this work.

## 6.1 Key contributions

The work presented in this thesis has contributed to the field of nonlinear aeroelasticity of very flexible aircraft. The focus has been mostly centered on understanding the low altitude dynamics of these extremely long endurance platforms. At the beginning of this effort, the individual solvers that could enable to solve these problem were close to mature enough. However, there was a need for a multidisciplinary integration exercise. The key contributions of this work are listed next.

- A versatile, efficient, high-quality nonlinear aeroelastic simulation tool has been designed, written and validated.
- Nonlinear time-domain simulations of slow very flexible aircraft have been conducted in turbulent fields generated with spectral (von Kármán and Kaimal) and high-fidelity numerical (LES) methods in order to assess the differences in predicted loads. It has been shown how higher turbulence intensities result in stronger discrete events, which cannot be simulated with spectral methods. Unidimensional turbulence generation methods severely underpredict the structural loads, overestimating the inertial relief effects.

- A complex nonlinear manoeuvre (a catapult-assisted take off) has been described, parametrised and simulated. This was chosen as a representative example of a highly unsteady nonlinear manoeuvre where linearisation is not possible. Results show that the link between acceleration, release velocity and ramp length is the most critical in the performance of the launch. The quicker the acceleration is, the further away from quasi-steady conditions the structure is at the release instant, which in turn, has a strong effect on the ground clearance.

- An original hybrid global/local optimisation based on Bayesian Optimisation and a gradient-less local method has been created with nonlinear aeroelastic problems in mind. It has been applied to the catapult-assisted take off problem, with a set of constraints and a cost function designed to account for ground clearance and wing root loads.

- An accurate aeroelasticity model of the X-HALE platform has been created in collaboration with University of Michigan. A verification exercise, covering the structural, aerodynamic and aeroelastic solvers has been conducted.

## 6.2 Outcomes

This thesis has dealt with the challenges associated with the computationally-efficient aeroelastic simulation of very flexible aircraft for loads and stability prediction. The developments presented here have consolidated and further expanded previous work performed by members of the Loads Control and Aeroelastics Lab, and past experience of the author working with a prototype implementation (SHARP) in an industrial setting.

While the current state of the art for nonlinear aeroelasticity of very flexible aircraft has come a long way in two decades, it has yet to reach a level of maturity that matches the requirements posed by current solar-powered HALE prototypes. The fact that a majority of the HALE-HAPS projects have failed or have been suspended indefinitely shows that it is not a minor enterprise, from an engineering and financial perspective. Part of this is due to the lack of simulation tools that are adequate for these configurations. Lower confidence in the analysis results in higher risks in the execution of the program (or suboptimal solutions), where design challenges need to be identified as early as possible. It is worth emphasising that of the full-sized prototypes that were flown, a large fraction of them have suffered loss of control and/or structural damage due to atmospheric disturbances. The majority of these being at low altitude.

### 6.2.1 A versatile, open-source nonlinear aeroelastic solver

Part of the work conducted by the author has been focused on designing a software architecture that can meet the demands of a highly multidisciplinary simulation, balancing the need for a versatile simulation and model input, while being simple enough to encourage new users to experiment with it. With this structure already set in place, a sustainable ecosystem of aeroelastic tools has been developed and integrated in SHARPy, not only for very flexible aircraft, but for high aspect-ratio wind turbines as well.

Other than refactoring existing algorithms, improvements in almost every part of the code have been included. Amongst most important ones are: the gravity loads implementation, where the accuracy reached before with 10 quadratic elements is now reached with 2, an improved vortex core model that improves the stability and accuracy of the UVLM, an added mass stabilisation scheme for use in coupled solutions, and a new mixed aeroelastic coupling scheme that allows to capture higher structural frequencies while maintaining geometrical nonlinearities in the aerodynamics solver. This allows the user to perform unique simulations that were not possible with the previous existing tools. Examples of this would be a fully coupled nonlinear aeroelastic simulations in closed loop while immersed in a three-dimensional, time dependent turbulent boundary layer, or a catapult-assisted take off manoeuvre, where a very flexible aircraft is accelerated from rest to close to trim flight while clamped at a few structural locations. One can go even further, and leveraging the modular nature of SHARPy, perform the catapult take off while immersed in the unsteady atmospheric boundary layer.

At the time this is written, SHARPy can be used for linear and nonlinear aeroelastic simulations of wings, aircraft and wind turbines both in open- and closed-loop, in arbitrary 3D flow fields. External controllers can be connected through a network protocol, reduced order models can be obtained and output, while the structural solver supports multibody constraints such as hinges, joints and prescribed nodal motions. The UVLM solver handles simulations more than double the size of what alternative codes can, while the structural solver is now two orders of magnitude faster than the original[1]. This is released under an Open-Source License and available free of charge. The code has been peer-reviewed when submitting to the Journal of Open Source Software [38], user documentation is automatically generated from the code, and every modification of the code is peer-reviewed and tests are run before being integrated in the source code. All the dependencies and related tools (such as compilers), are freely available.

[1]Not all of these contributions are due to the author. SHARPy now a collaborative effort.

### 6.2.2 Low altitude atmospheric disturbance effect on very flexible aircraft

This thesis has carried out a pioneer study on physics-generated background flows interaction with very flexible aircraft. The aeroelastic coupling and dynamic loading on a very flexible aircraft has been found to be highly modulated by the coherent structures and flow features of low-altitude atmospheric turbulence. This may be of critical importance for current and future HAPS prototypes.

The choice of the turbulence generation method has been identified as a significant variable on both loads and flight dynamics response of the aircraft. Of the three models studied in this work, the ones with at least 2-D coherence, LES and Kaimal, provide similar results, although LES was able to predict greater peak loads caused by discrete events, which provide intermittence, one of the properties of real life turbulence that cannot be directly replicated with spectral methods. Kaimal was found to be problematic when it is used within the moving frame of reference as the application of the frozen turbulence hypothesis used to transform temporal 2D planes into spatial full three-dimensional wind fields does not generate three-dimensional coherence in the wind field. As a result, small deviations of the aircraft from the frozen turbulence convection velocity generated sudden variations in frequency content of the flow which affected the load envelopes in our simulations. LES, on the other hand, simulates a more realistic turbulent environment with large scales extending in the streamwise direction and generating flow coherence. The von Kármán model was the worst performing method of those considered, particularly when used in a moving three-dimensional frame of reference. The uniform planar velocity fields from this approach result in spanwise loading which overstates the inertial relief importance, underpredicting the magnitude of the loads and overpredicting the energy transferred into rigid body motion. Therefore its use for loads and aeroelastics at low altitudes is not advised as the only continuous loads prediction method in the high aspect-ratio, very flexible aircraft context.

The comparison between the flexible and the stiff models has demonstrated the extent of the aeroelastic effects in the previous simulations. Analysis of these two aircraft flight dynamics shows that stiffer aircraft are prone to instabilities from high-frequency perturbations leading to large out-of-plane moment loads. More flexible vehicles, on the other hand, have natural load alleviation mechanisms through structural deformation, and rigid body motion coupling. Results have shown smaller out-of-plane loads for the flexible configuration, as the energy contained in gust loading is distributed between

rigid body motion, torsion, OOP and IP moments. Overall, as the flight dynamics of high-aspect ratio flexible aircraft were found to be fundamentally different from that of the stiff one, more emphasis should be given to non-linear aeroelastic phenomena which might occur under turbulent conditions and are currently not well-captured by current standard methods.

Results from Chapter 4 have illustrated the need for continuous loads criteria other than the classic von Kármán spectrum for HAPS-like configurations. The typical low speed, high aspect-ratio HALE aircraft are highly sensitive to structural cross-couplings that might not be evident when exciting the structure with span-constant disturbances. These can, at the same time, overstate the influence of the inertia relief, underpredicting the resulting loads. Lastly, discrete intermittent events can be the ones causing maximum loads, which cannot be simulated with spectral methods. In conclusion, the author suggests that turbulence generation methods with at least 2-D coherence should be considered when dealing with the low altitude stage of the flight of HALE-like platforms.

### 6.2.3 Catapult-Assisted Take Off simulation

A nonlinear aeroelastic simulation strategy has been presented in Section 6.2.1. It allows to simulate complex maneuvers of very flexible aircraft.

Catapult-assisted takeoff has been used as a representative example of a highly dynamical nonlinear aeroelastic problem. In this problem, the aircraft is attached to the catapult on multiple nodes and the velocity of these nodes is prescribed so that a trajectory and acceleration law is enforced. Once the release speed is reached, the aircraft is released and left to fly in open loop. This problem features a high range of flight speeds (from zero to close to cruise speed), high range of structural deformations (from gravity-only loading to aeroelastic equilibrium in cruise) and changing driving mechanisms as the simulation progresses. Due to this, linear aeroelastic simulation methods would not be able to capture the complex phenomena that arise from the interaction of structure, aerodynamics, prescribed rigid-body motion and flight dynamics. To my knowledge, no comparable simulation has ever been reported in the open literature on aeroelasticity.

The proposed CATO procedure is described using four parameters, namely, the acceleration value, the release speed, the ramp angle and the offset attitude of the aircraft. The ramp is also assumed to be straight, the pitch angle of the aircraft constant and the acceleration rate constant until release.

A baseline case has been defined and simulated for the very flexible model and a very stiff version with identical mass properties and geometry. This comparison allows to quantify the influence of flexibility effects in the CATO maneuver. The results show a very different behaviour in structural loads and post-release flight dynamics between the two models. From a loads point of view, the time-varying effective dihedral in the flexible case induced a strong coupling between torsion and in-plane bending. Peak-to-peak torsional and in-plane bending loads are more than doubled in the case of the flexible model with respect to the stiff one.

A parametric investigation has shown that ground clearance is very sensitive to release speed and pitch angle. Wing root loads are increased by large acceleration values while a high release speed reduces them. Longer ramps are therefore preferred, and a ramp length constraint need to be considered. Depending on the acceleration and release speed combination, the characteristic time of the CATO is from 60 to 108% of the period of the first bending mode of the wing. This, together with the large damping of out-of-plane motions, explains why larger ramps with higher release speed perform well also from a ground clearance point of view. In limited length ramps, the CATO is too fast for the wing to adopt the quasi-steady deflection for those instantaneous conditions. On release, the centre of mass, which is located close to the wing root due to the payload, plunges due to the insufficient lift generated by the wings, which then deform and achieve a deflection close to the corresponding steady one. This plunging and recovering motion can excite rigid body modes such as the phugoid, and depending on the controllability of the aircraft and the ground clearance it might be a limiting factor to consider when designing the CATO procedure.

### 6.2.4 Optimisation of nonlinear maneuvers

The efficiency of SHARPy as a nonlinear aeroelastic solver allows, for the first time in the literature, to optimise a full very flexible aircraft complex manoeuvre. Given the complexity of the CATO manoeuvre, the definition of an "optimal" launch is far from a straightforward. This is usually the case with practical optimisation problems. This is due to a number of reasons, the most important ones being: expensive cost function evaluation, lack of Jacobian/Hessian approximations or analytical expression for them, potentially complex constraints and cost function shape. The lack of derivative information, together with the expensive cost function evaluation, makes any gradient-based optimisation algorithm not useful for our purposes. Another option, quite well

known, are genetic algorithms. These are versatile and the literature around them is extensive. However, they require a great number of function calls, do not build a surrogate for the later local optimisation stage, and a certain degree of fine tuning of the parameters, which vary from problem to problem, is needed.

This work proposes a hybrid global/local optimisation strategy developed by the author. The global stage of the optimisation is based on a Bayesian Optimisation framework, a technique based on a building a surrogate of the cost function assuming it can be represented as a combination of Gaussian processes (Gaussian Process Regression). Sampling points are decided using an acquisition function, which tries to balance the minimisation of the value of the cost function while minimising the uncertainty of the surrogate in the areas of interest. This process can be understood as a global smart sampling algorithm, where low cost and high variability regions are sampled densely. Once the global optimisation has identified the global minimum, the surrogate built in the previous step is fed into a local gradient-less Powell method to refine the solution. Once it is found, the surrogate approximation is checked against the real nonlinear aeroelastic simulation. When the discrepancy is small, the minimisation problem is complete, while if not, the Gaussian Process Regression model is updated with the new sample and the local optimiser is run again. The cases shown in this work have been converged in less than a day running 4 simulations in parallel, resulting in around 160 SHARPy calls.

Bayesian optimization has been shown to provide a suitable strategy to define dynamic maneuvers that can be parametrised with relatively small number of design variables. The resulting maneuver, which is defined under ideal conditions (e.g., calm air) can then be used as the tracked reference for a closed-loop control scheme.

Specifically for the CATO problem, a cost function based on wing root loads and ground clearance has been defined. In addition, a set of constraints defining the acceptable domain have been declared. One of them restricts the ramp length through a combination of acceleration and release velocity. The other limits the effective pitch angle of the aircraft in order to avoid operating the aircraft anywhere close to the stall regime.

The optimization studies have shown that, for a given model and constraints, there is a minimal ramp length under which the aircraft would hit the ground. In addition, it has been seen that the parameter with the strongest effect in the solution is the release speed. Unsurprisingly, with greater release velocities the ground clearance is larger and the wing root loads are smaller, as long as the ramp length constraints allow the acceleration to be small enough. More interestingly however, the ramp angle parameter has a very small effect on the launch, unlike the pitch angle parameter. A large pitch angle parameter

encourages wing loading during the early stages of the CATO, which results in a closer to quasi-steady structural state at release. However, a large incidence angle also increases the wing root loads. In our results the incidence angle is not affected by the optimization constraints, as the increase of loads balances the ground clearance effect.

## 6.3 Proposals for future work

The aeroelastic framework presented in this work has proven to be highly versatile in the kind of aeroelasticity problems it can tackle. The cases included in this thesis are a good representation of it. This enables a broad range of possibilities concerning future developments. However, there is room for further work in a number of areas directly related to this thesis. These will be discussed next, together with some developments that would complement the current work.

- **The aerodynamic solver**, now based on an unsteady vortex-lattice method, is the most expensive part of the simulation. Acceleration methods such as the Fast Multipole Method, have shown potential to greatly reduce the cost of the aerodynamic simulations, specially the ones with large wakes that cannot be simply truncated (e.g. a wind turbine).

- **Validation with real flight data** would be a very important step in the field. To the knowledge of the author, no comparison of numerical results with flight data has been conducted for very flexible aircraft. The lack of good quality flight data is the main constraint.

- **Reduced order models** that can run in real-time have the potential of turning the code into a full-fledged flight simulator physics engine, including complex dynamics that are not available in any commercially-available simulation. In manned (or remotely piloted) vehicles with large flexibility, the importance of receiving feedback from the pilot before any prototype is finished, together with having an accurate simulator of the aircraft, cannot be understated.

- **Higher-fidelity CFD solvers** could be used instead of potential methods for certain simulations, where nonlinear aerodynamic effects are important, such as high-lift configurations, or stalled flight conditions. These also offer the possibility of improving the accuracy of the drag prediction, with the inclusion of viscous effects.

- **The unsteadiness of the turbulent boundary layer** increases the cost and complexity of the simulations, increasing memory consumption and CPU time. In conventional aircraft, the speed of the aircraft justifies the assumption for steady flow fields. The question remains open for low-speed, large span platforms.

- **The robustness of the optimal CATO solutions** found for calm atmosphere needs to be studied and quantified once the incoming flow is different from the one used during the optimisation. The differences can include discrepancies in wind speed, direction, or even turbulence intensity (this would effectively link the two main simulations presented in this work).

- **Closed-loop manoeuvres** are a necessary simulation to understand first the influence of the control loop in the CATO procedure and second, if the performance of the manoeuvre can be improved, from a ground clearance or loads standpoint. Furthermore, the control law could be optimised simultaneously with the launch parameters.

# References


[1] Adeli, H., Weaver, W., and Gere, J. M. (1978). Algorithms for nonlinear structural dynamics. *Journal of the Structural Division*, 104(2).

[2] Afonso, F., Vale, J., Oliveira, É., Lau, F., and Suleman, A. (2017). A review on non-linear aeroelasticity of high aspect-ratio wings. *Progress in Aerospace Sciences*, 89(2016):40–57.

[3] Agrawal, O. P. and Shabana, A. A. (1985). Dynamic analysis of multibody systems using component modes. *Computers & Structures*, 21(6):1303–1312.

[4] Agrawal, O. P. and Shabana, A. A. (1986). Application of deformable-body mean axis to flexible multibody system dynamics. *Computer Methods in Applied Mechanics and Engineering*, 56(2):217–245.

[5] Ahmadi, A. R. and Widnall, S. E. (1985). Unsteady lifting-line theory as a singular-perturbation problem. *Journal of Fluid Mechanics*, 153:59–81.

[6] Albano, E. (1969). A doublet-lattice method for calculating lift distributions on oscillating surfaces in subsonic flows. *AIAA Journal*, 7(2):279–285.

[7] Artola, M., Goizueta, N., Wynn, A., and Palacios, R. (2020). Modal-based nonlinear estimation and control for highly flexible aeroelastic systems. In *AIAA Scitech 2020 Forum*. American Institute of Aeronautics and Astronautics.

[8] Bao, Y., Palacios, R., Graham, M., and Sherwin, S. (2016). Generalized thick strip modelling for vortex-induced vibration of long flexible cylinders. *Journal of Computational Physics*, 321:1079–1097.

[9] Barbosa, G. C., Bertolin, R., González, P. J., Guimaraes Neto, A. B., and Silvestre, F. J. (2018). Fuzzy gain-scheduling applied for a very flexible aircraft. In *2018 AIAA Guidance, Navigation, and Control Conference*.

[10] Bauchau, O. A. (2011). *Flexible Multibody Dynamics*, volume 176 of *Solid Mechanics and Its Applications*. Springer Netherlands, Dordrecht.

[11] Bertin, J. J. and Smith, M. L. (1998). *Aerodynamics for engineers*, volume 5. Prentice Hall Upper Saddle River, NJ.

[12] Bisplinghoff, R. L., Ashley, H., and Halfman, R. L. (2013). *Aeroelasticity*. Courier Corporation.

[13] Blair, M. (1992). A Compilation of the Mathematics Leading to the Doublet Lattice Method. Technical report, Wright Lab Wright-Patterson AFB OH.

[14] Boschnitsch, A., William Usab, J., and Epstein, R. (1999). Fast lifting panel method. In *14th Computational Fluid Dynamics Conference*, number c, pages 99–3376.

[15] Bou-Zeid, E., Meneveau, C., and Parlange, M. (2005). A scale-dependent lagrangian dynamic model for large eddy simulation of complex turbulent flows. *Physics of Fluids*, 17(2):025105.

[16] Brase, L. O. and Eversman, W. (1988). Application of transient aerodynamics to the structural nonlinear flutter problem. *Journal of Aircraft*, 25(11):1060–1068.

[17] Campbell, C. W. (1986). Monte Carlo Turbulence Simulation Using Rational Approximations to von Kármán Spectra. *AIAA Journal*, 24(1):62–66.

[18] Canavin, J. R. and Likinst, P. W. (1977). Floating Reference Frames for Flexible Spacecraft. *Journal of Spacecraft and Rockets*, 14(12).

[19] Cardona, A. (2000). Superelements Modelling in Flexible Multibody Dynamics. *Multibody System Dynamics*, 4(2/3):245–266.

[20] Cardona, A. and Geradin, M. (1988). A beam finite element non-linear theory with finite rotations. *International Journal for Numerical Methods in Engineering*, 26(11):2403–2438.

[21] Cardona, A. and Geradin, M. (1991). Modelling of superelements in mechanism analysis. *International Journal for Numerical Methods in Engineering*, 32(8):1565–1593.

[22] Carson Yates Jr., E. (1987). AGARD standard aeroelastic configurations for dynamic response. Candidate configuration I.-wing 445.6. Technical report.

[23] Cavin III, R. K. and Dusto, A. R. (1977). Hamilton's Principle: Finite-Element Methods and Flexible Body Dynamics. *AIAA Journal*, 15(12):1684–1690.

[24] Cesnik, C. E. S. and Brown, E. (2002a). Modeling of High Aspect Ratio Active Flexible Wings for Roll Control. In *43rd AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference*, Reston, Virigina. American Institute of Aeronautics and Astronautics.

[25] Cesnik, C. E. S. and Brown, E. L. (2002b). Modelling of high aspect ratio active flexible wings for roll control. In *43rd Structures, Structural Dynamics, and Material Conference*, Denver, Colorado.

[26] Cesnik, C. E. S. and Hodges, D. H. (1997). VABS: A New Concept for Composite Rotor Blade Cross-Sectional Modeling. *Journal of the American Helicopter Society*, 42(1):27–38.

[27] Cestino, E. (2006). Design of solar high altitude long endurance aircraft for multi payload and operations. *Aerospace Science and Technology Cestino / Aerospace Science and Technology*, 10(10):541–550.

[28] Chabalko, C., Fitzgerald, T., and Balachandran, B. (2013). GPGPU implementation and benchmarking of the unsteady vortex lattice method. In *51st AIAA Aerospace Sciences Meeting including the New Horizons Forum and Aerospace Exposition*, Reston, Virigina. American Institute of Aeronautics and Astronautics.

[29] Cook, R. G., Palacios, R., and Goulart, P. (2013). Robust gust alleviation and stabilization of very flexible aircraft. *AIAA Journal*, 51(2).

[30] Cook, R. G., Wales, C. J., Gaitonde, A. L., Jones, D. P., and Cooper, J. E. (2018). Uncertainty quantification in gust loads analysis of a highly flexible aircraft wing. In *International Forum on Aeroelasticity and Structural Dynamics, Como, Italy*.

[31] Cooper, J. E., Cook, R. G., Francois, G., Torre Rodríguez, Ó. d. l., Neild, S. A., Lowenberg, M. H., Alexander, S., Coetzee, E., and Evans, M. (2019). Wind tunnel testing of a high aspect ratio wing model. In *Proceedings of Forum of Aeroelasticity and Structural Dynamics IFASD 2019, 9-13 June 2019, Savannah, Georgia, USA*, pages 094–1. Curran.

[32] Cordero-Gracia, M., Gómez, M., Ponsin, J., and Valero, E. (2012). An interpolation tool for aerodynamic mesh deformation problems based on octree decomposition. In *Aerospace Science and Technology*.

[33] Craig, R. R. and Bampton, M. C. C. (1968). Coupling of substructures for dynamics analyses. *AIAA Journal*, 6(7):1313.

[34] Crisfield, M. A. and Jelenić;, G. (1999). Objectivity of strain measures in the geometrically exact three-dimensional beam theory and its finite-element implementation. *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, 455(1983).

[35] Danielson, D. A. and Hodges, D. H. (1987). Nonlinear Beam Kinematics by Decomposition of the Rotation Tensor. *Journal of Applied Mechanics*, 54(2):258.

[36] Davoudi, B., Taheri, E., Duraisamy, K., Jayaraman, B., and Kolmanovsky, I. (2019). Quad-rotor flight simulation in realistic atmospheric conditions. *arXiv preprint arXiv:1902.01465*.

[37] de C. Henshaw, M., Badcock, K., Vio, G., Allen, C., Chamberlain, J., Kaynes, I., Dimitriadis, G., Cooper, J., Woodgate, M., Rampurawala, A., Jones, D., Fenwick, C., Gaitonde, A., Taylor, N., Amor, D., Eccles, T., and Denley, C. (2007). Non-linear aeroelastic prediction for aircraft applications. *Progress in Aerospace Sciences*, 43(4):65 – 137.

[38] del Carre, A., Muñoz, A., Goizueta, N., and Palacios, R. (2019). SHARPy: A Dynamic Aeroelastic Simulation Toolbox for Very Flexible Aircraft and Wind Turbines. *Journal of Open Source Software*, 4(44).

[39] del Carre, A., Muñoz-Simón, A., Goizueta, N., and Palacios, R. (2019). SHARPy: A dynamic aeroelastic simulation toolbox for very flexible aircraft and wind turbines. *Journal of Open Source Software*, 4(44):1885.

[40] Del Carre, A., Teixeira, P., Cesnik, C. E. S., and Palacios, R. (2019). Nonlinear response of a very flexible aircraft under lateral gust. In *IFASD 2019 Forum*, Savannah, Georgia.

[41] Demasi, L. and Livne, E. (2009). Dynamic Aeroelasticity of Structurally Nonlinear Configurations Using Linear Modally Reduced Aerodynamic Generalized Forces. *AIAA Journal*, 47(1):70–90.

[42] Deskos, G., Laizet, S., and Palacios, R. (2020). Winc3d: A novel framework for turbulence-resolving simulations of wind farm wake interactions. *Wind Energy*, 23(3):779–794.

[43] Deskos, G., Laizet, S., and Piggott, M. D. (2018). Development and validation of the higher-order finite-difference wind farm simulator, WInc3D. In *3rd International Conference on Renewable Energies Offshore (RENEW2018)*, Lisbon, Portugal, 8–10 October 2018.

[44] Diederich, F. and Drischler, J. (1957). Effect of spanwise variations in gust intensity on the lift due to atmospheric turbulence. national advisory committee for aeronautics. Technical report, Technical Note 3920.

[45] Dizy, J., Palacios, R., and Pinho, S. T. (2013). Homogenisation of slender periodic composite structures. *International Journal of Solids and Structures*, 50(9):1473–1481.

[46] D'Oliveira, F. A., de Melo, F. C. L., and Devezas, T. C. (2016). High-altitude platforms - present situation and technology trends. *Journal of Aerospace Technology and Management*, 8(3):249–262.

[47] Drela, M. (1999). Integrated Simulation Model for Preliminary Aerodynamic, Structural, and Control-Law Design of Aircraft. *40th AIAA Structures, Structural Dynamics and Materials Conference*, pages 1644–1656.

[48] EASA (2007). Certification Specifications for Large Aeroplanes (CS-25).

[49] Eichenbaum, F. D. (1971). A general theory of aircraft response to three-dimensional turbulence. *Journal of Aircraft*, 8(5):353–360.

[50] Etkin, B. (1982). *Dynamics of flight: stability and control*. John Wiley & Sons Australia, Limited.

[51] Etkin, B. (2012). *Dynamics of atmospheric flight*. Courier Corporation.

[52] European Aviation Safety Agency (2020a). Certification specifications and acceptable means of compliance for large aeroplanes cs-25. amendment 25. Technical report, European Aviation Safety Agency.

[53] European Aviation Safety Agency (2020b). Certification specifications for normal-category aeroplanes (cs-23) and acceptable means of compliance and guidance material to the certification specifications for normal-category aeroplanes(amc & gm to cs-23). Technical report, European Aviation Safety Agency.

[54] Felippa, C. A., Park, K., and Ross, M. (2011). A classification of interface treatments for fsi. In *Fluid Structure Interaction II*, pages 27–51. Springer.

[55] Frazier, P. I. (2018). A tutorial on bayesian optimization. *arXiv preprint arXiv:1807.02811*.

[56] Friedmann, P., McNamara, J., Thuruthimattam, B., and Nydick, I. (2004). Aeroelastic analysis of hypersonic vehicles. *Journal of Fluids and Structures*, 19(5):681–712.

[57] Fung, Y. C. (2008). *An introduction to the theory of aeroelasticity*. Courier Dover Publications.

[58] Gaitonde, D. V. and Visbal, M. R. (1998). High-order schemes for navier-stokes equations: Algorithm and implementation into fdl3di. Technical report, Air Force Research Laboratory.

[59] Garcia, J. and Guruswamy, G. (1999). Aeroelastic analysis of transonic wings using navier-stokes equations and a nonlinear beam finite element model. In *40th Structures, Structural Dynamics, and Materials Conference and Exhibit*. American Institute of Aeronautics and Astronautics.

[60] García, J., Muñoz-Paniagua, J., and Crespo, A. (2017). Numerical study of the aerodynamics of a full scale train under turbulent wind conditions, including surface roughness effects. *Journal of Fluids and Structures*, 74:1–18.

[61] Garrick, I. E. (1937). Propulsion of a flapping and oscillating airfoil.

[62] Géradin, M. and Cardona, A. (2001). *Flexible multibody dynamics: a finite element approach*. Wiley.

[63] Géradin, M. and Rixen, D. J. (2014). *Mechanical vibrations: theory and application to structural dynamics*. John Wiley & Sons.

[64] Geuzaine, P., Brown, G., Harris, C., and Farhat, C. (2003). Aeroelastic Dynamic Analysis of a Full F-16 Configuration for Various Flight Conditions. *AIAA Journal*, 41(3):363–371.

[65] Gonzalo, J., López, D., Domínguez, D., García, A., and Escapa, A. (2018). On the capabilities and limitations of high altitude pseudo-satellites. *Progress in Aerospace Sciences*, 98.

[66] Haghighat, S., Martins, J. R. R. A., and Liu, H. H. T. (2012). Aeroservoelastic Design Optimization of a Flexible Wing. *Journal of Aircraft*, 49(2):432–443.

[67] Hallissy, B. and Cesnik, C. E. S. (2011). High-fidelity Aeroelastic Analysis of Very Flexible Aircraft. In *52nd AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics and Materials Conference*, Reston, Virigina. American Institute of Aeronautics and Astronautics.

[68] Haugen, D. A., Kaimal, J. C., and Bradley, E. F. (1971). An experimental study of reynolds stress and heat flux in the atmospheric surface layer. *Quarterly Journal of the Royal Meteorological Society*, 97(412):168–180.

[69] Hesse, H. (2013). *Consistent aeroelastic linearisation and reduced-order modelling in the dynamics of manoeuvring flexible aircraft*. PhD thesis, Imperial College London.

[70] Hesse, H., Palacios, R., and Murua, J. (2014). Consistent Structural Linearization in Flexible Aircraft Dynamics with Large Rigid-Body Motion. *AIAA Journal*, 52(3).

[71] Hilber, H. M., Hughes, T. J., and Taylor, R. L. (1977). Improved numerical dissipation for time integration algorithms in structural dynamics. *Earthquake Engineering & Structural Dynamics*, 5(3):283–292.

[72] Hoblit, F. M. (1988). *Gust Loads on Aircraft: Concepts and Applications*. American Institute of Aeronautics and Astronautics.

[73] Hodges, D. H. (1990). A mixed variational formulation based on exact intrinsic equations for dynamics of moving beams. *International Journal of Solids and Structures*, 26(11):1253–1273.

[74] Hodges, D. H. (2003). Geometrically Exact, Intrinsic Theory for Dynamics of Curved and Twisted Anisotropic Beams. *AIAA JOURNAL*, 41(6).

[75] Hodges, D. H. (2006). Nolinear composite beam theory. *Progress in astronautics and aeronautics*, 213:304.

[76] Houbolt, J. C. (1973). Atmospheric Turbulence. *AIAA Journal*, 11(4):421–437.

[77] Houghton, E. L. and Carpenter, P. W. (2003). *Aerodynamics for engineering students*. Elsevier.

[78] Hsiao, K. M., Lin, J. Y., and Lin, W. Y. (1999). A consistent co-rotational finite element formulation for geometrically nonlinear dynamic analysis of 3-D beams. *Computer Methods in Applied Mechanics and Engineering*, 169(1-2):1–18.

[79] Hunsaker, J. C. and Wilson, E. B. (1917). Report on behavior of aeroplanes in gusts. Technical report, NACA Technical Report 1, NACA-TR-1.

[80] Ibrahimbegović, A. (1995). On finite element implementation of geometrically nonlinear Reissner's beam theory: three-dimensional curved beam elements. *Computer Methods in Applied Mechanics and Engineering*, 122(1-2):11–26.

[81] International Electrotechnical Commission (2005). *IEC-61400-1 Wind turbines-Part 1: Design requirements*.

[82] Jeong, M.-S., Kim, S.-W., Lee, I., and Yoo, S.-J. (2014). Wake impacts on aerodynamic and aeroelastic behaviors of a horizontal axis wind turbine blade for sheared and turbulent flow conditions. *Journal of Fluids and Structures*, 50:66–78.

[83] Jones, E., Oliphant, T., Peterson, P., et al. (2001–). SciPy: Open source scientific tools for Python.

[84] Jones, J. R. (2017). *Development of a Very Flexible Testbed Aircraft for the Validation of Nonlinear Aeroelastic Codes*. PhD thesis, University of Michigan.

[85] Jones, J. R. and Cesnik, C. E. S. (2015). Preliminary flight test correlations of the X-HALE aeroelastic experiment. *The Aeronautical Journal*, 119(1217):855–870.

[86] Jones, R. T. (1940). The unsteady lift of a wing of finite aspect ratio. Technical report, National Advisory Committee for Aeronautics. NACA-TR-681.

[87] Jonkman, B. J. and Kilcher, L. (2012). Turbsim user's guide: Version 1.06.00. Technical report, National Renewable Energy Laboratory.

[88] Kaimal, J. C., Wyngaard, J. C., Izumi, Y., and Coté, O. R. (1972). Spectral characteristics of surface-layer turbulence. *Quarterly Journal of the Royal Meteorological Society*, 98(417):563–589.

[89] Katz, J. and Plotkin, A. (2001). *Low-Speed Aerodynamics*. Cambridge Aerospace Series. Cambridge University Press.

[90] Kennedy, G. J. and Martins, J. R. (2012). A homogenization-based theory for anisotropic beams with accurate through-section stress and strain prediction. *International Journal of Solids and Structures*, 49(1):54–72.

[91] Khosravi, P., Ganesan, R., and Sedaghati, R. (2007). Corotational non-linear analysis of thin plates and shells using a new shell element. *International Journal for Numerical Methods in Engineering*, 69(4):859–885.

[92] Kier, T. (2005). Comparison of Unsteady Aerodynamic Modelling Methodologies with Respect to Flight Loads Analysis. *AIAA Atmospheric Flight Mechanics Conference and Exhibit*, pages 1–14.

[93] Kier, T., Looye, G., Scharpenberg, M., and Reijerkerk, M. (2007). Process, methods and tools for flexible aircraft flight dynamics model integration. In *International Forum on Aeroelasticity and Structural Dynamics*.

[94] Konstadinopoulos, P., Thrasher, D., Mook, D., Nayfeh, A., and Watson, L. (1985). A vortex-lattice method for general, unsteady aerodynamics. *Journal of aircraft*, 22(1):43–49.

[95] Krenk, S. (2006). Energy conservation in newmark based time integration algorithms. *Computer Methods in Applied Mechanics and Engineering*, 195(44-47):6110–6124.

[96] Kuhl, D. and Crisfield, M. A. (1999). Energy-conserving and decaying algorithms in non-linear structural dynamics. *International Journal for Numerical Methods in Engineering*, 45(5):569–599.

[97] Laizet, S. and Lamballais, E. (2009). High-order compact schemes for incompressible flows: A simple and efficient method with quasi-spectral accuracy. *Journal of Computational Physics*, 228(16):5989 – 6015.

[98] Laizet, S. and Li, N. (2011). Incompact3d: A powerful tool to tackle turbulence problems with up to o($10^5$) computational cores. *International Journal for Numerical Methods in Fluids*, 67(11):1735–1757.

[99] Lens, E. V., Cardona, A., and Geradin, M. (2004). Energy preserving time integration for constrained multibody systems. *Multibody System Dynamics*, 11(1):41–61.

[100] Liepmann, H. W. (1952). On the application of statistical concepts to the buffeting problem. *Journal of the Aeronautical Sciences*, 19(12):793–800.

[101] Lim, J. (1990). *Two-dimensional signal and image processing*. Prentice Hall, Englewood Cliffs, N.J.

[102] Liu, Y., Xie, C., Yang, C., and Cheng, J. (2016). Gust response analysis and wind tunnel test for a high-aspect ratio wing. *Chinese Journal of Aeronautics*, 29(1):91 – 103.

[103] Lucia, D. J., Beran, P. S., and Silva, W. A. (2004). Reduced-order modeling: new approaches for computational physics. *Progress in Aerospace Sciences*, 40(1):51 – 117.

[104] Ly, U. and Chan, Y. (1980). Time-domain computation of aircraft gust covariance matrices. In *6th AIAA Atmospheric Flight Mechanics Conference*, Danvers, Massachussetts, USA. AIAA Paper No 1980-1615.

[105] Mann, J. (1994). The spatial structure of neutral atmospheric surface-layer turbulence. *Journal of Fluid Mechanics*, 273:141–168.

[106] Mann, J. (1998). Wind field simulation. *Probabilistic Engineering Mechanics*, 13(4):269 – 282.

[107] Maraniello, S. and Palacios, R. (2017). Optimal rolling maneuvers with very flexible wings. *AIAA Journal*, 55(9).

[108] Margason, R. J., Kjelgaard, S. O., Sellers III, W. L., Morris Jr., C. E. K., Walkley, K. B., and Shields, E. W. (1985). Subsonic panels methods - A comparison of several production codes. In *23rd Aerospace Sciences Meeting*.

[109] Mason, P. J. and Thomson, D. J. (1992). Stochastic backscatter in large-eddy simulations of boundary layers. *Journal of Fluid Mechanics*, 242:51–78.

[110] Maute, K., Nikbay, M., and Farhat, C. (2001). Coupled Analytical Sensitivity Analysis and Optimization of Three-Dimensional Nonlinear Aeroelastic Systems. *AIAA Journal*, 39(11):2051–2061.

[111] Meirovitch, L. (1989). State Equations of Motion for Flexible Bodies in Terms of Quasi-Coordinates. *IFAC Symposium on Dynamics of Controlled Mechanical Systems*, 1(5):37 – 48.

[112] Meirovitch, L. and Tuzcu, I. (2007). The Lure of the Mean Axes. *Journal of Applied Mechanics*, 74(3):497.

[113] Melin, T. (2000). A vortex lattice MATLAB implementation for linear aerodynamic wing applications. *Master's Thesis, Department of Aeronautics, Royal Institute of Technology (KTH), Stockholm, Sweden*.

[114] Melin, T. (2013). Tornado. *KTH, Stockholm, Sweden, Masters Thesis and continued development. http://www.flyg.kth.se/divisions/aero/software/tornado*.

[115] MIL-HDBK-1797 (1997). U.S. Military Handbook MIL-HDBK-1797. Technical report, Department of Defense, USA.

[116] Miles, J. W. (1956). The aerodynamic force on an airfoil in a moving gust. *Journal of the Aeronautical Sciences*, 23(11):1044–1050.

[117] Milne, R. D. (1962). Dynamics of the deformable plane. *Her Majesty's Stationary Office TR-3345, London, England, UK*.

[118] Milne, R. D. (1968). Some remarks on the dynamics of deformable bodies. *AAIA Journal*, 6(3):556–558.

[119] Moeng, C.-H. (1984). A large-eddy-simulation model for the study of planetary boundary-layer turbulence. *Journal of the Atmospheric Sciences*, 41(13):2052–2062.

[120] Moler, C. and Loan, C. V. (2003). Nineteen dubious ways to compute the exponential of a matrix, twenty-five years later. *SIAM Review*, 45(1):3–49.

[121] Monin, A. S. and Obukhov, A. (1954). Basic laws of turbulent mixing in the surface layer of the atmosphere. *Contrib. Geophys. Inst. Acad. Sci. USSR*, 151:163–187. Monin1954.

[122] Muñoz-Simón, A., Wynn, A., and Palacios, R. (2020). Unsteady and three-dimensional aerodynamic effects on wind turbine rotor loads. In *AIAA Scitech 2020 Forum*. American Institute of Aeronautics and Astronautics.

[123] Murua, J. (2012). *Flexible Aircraft Dynamics with a Geometrically-Nonlinear Description of the Unsteady Aerodynamics*. PhD thesis, Imperial College London.

[124] Murua, J., Palacios, R., and Graham, J. M. R. (2010). Modeling of Nonlinear Flexible Aircraft Dynamics Including Free-Wake Effects. *AIAA Atmospheric Flight Mechanics*, (August):1–25.

[125] Murua, J., Palacios, R., and Graham, J. M. R. (2012a). Applications of the unsteady vortex-lattice method in aircraft aeroelasticity and flight dynamics. *Progress in Aerospace Sciences*, 55.

[126] Murua, J., Palacios, R., and R. Graham, J. M. (2012b). Assessment of wake-tail interference effects on the dynamics of flexible aircraft. *AIAA Journal*, 50(7).

[127] NATO (2019). Standarized agreement 4671. AUV airworthiness requirements. Technical report, NATO.

[128] Nelder, J. A. and Mead, R. (1965). A simplex method for function minimization. *The computer journal*, 7(4).

[129] Newmark, N. M. et al. (1959). A method of computation for structural dynamics. American Society of Civil Engineers.

[130] Noll, T. E., Brown, J. M., Perez-davis, M. E., Ishmael, S. D., Tiffany, G. C., and Gaier, M. (2004). Investigation of the Helios Prototype Aircraft Mishap. Technical Report January, NASA.

[131] Offner, G., Eizenberger, T., and Priebsch, H. H. (2006). Separation of reference motions and elastic deformations in an elastic multi-body system. *Proceedings of the Institution of Mechanical Engineers, Part K: Journal of Multi-body Dynamics*, 220(1):63–75.

[132] Palacios, R. and Cesnik, C. E. S. (2005). Static nonlinear aeroelasticity of flexible slender wings in compressible flow. In *46th AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics and Materials Conference*. American Institute of Aeronautics and Astronautics.

[133] Palacios, R., Murua, J., and Cook, R. (2010). Structural and aerodynamic models in nonlinear flight dynamics of very flexible aircraft. *AIAA Journal*, 48(11):2648–2659.

[134] Patil, M. and Hodges, D. (2004). On the importance of aerodynamic and structural geometrical nonlinearities in aeroelastic behavior of high-aspect-ratio wings. *Journal of Fluids and Structures*, 19(7):905–915.

[135] Patil, M. J., Hodges, D., and Cesnik, C. E. S. (1998). Nonlinear aeroelastic analysis of aircraft with high-aspect-ratio wings. In *39th AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference and Exhibit*. American Institute of Aeronautics and Astronautics.

[136] Patil, M. J., Hodges, D. H., and Cesnik, C. E. S. (2000). Nonlinear aeroelastic analysis of complete aircraft in subsonic flow. *Journal of Aircraft*, 37(5):753–760.

[137] Patil, M. J., Hodges, D. H., and Cesnik, C. E. S. (2001a). Limit-Cycle Oscillations of High-Aspect-Ratio Wings. *Journal of Fluids and Structures*, 15(1):107–132.

[138] Patil, M. J., Hodges, D. H., and Cesnik, C. E. S. (2001b). Nonlinear aeroelasticity and flight dynamics of high-altitude long-endurance aircraft. *Journal of Aircraft*, 38(1):88–94.

[139] Phillips, W. F. (2010). *Mechanics of flight*. John Wiley & Sons, Inc, Hoboken, New Jersey.

[140] Pope, S. (2000). *Turbulent Flows*. Cambridge University Press.

[141] Porté-Agel, F., Meneveau, C., and Parlange, M. B. (2000). A scale-dependent dynamic model for large-eddy simulation: application to a neutral atmospheric boundary layer. *Journal of Fluid Mechanics*, 415:261–284.

[142] Powell, M. J. (1964). An efficient method for finding the minimum of a function of several variables without calculating derivatives. *The computer journal*, 7(2).

[143] Qi, P., Zhao, X., Wang, Y., Palacios, R., and Wynn, A. (2018). Aeroelastic and trajectory control of high altitude long endurance aircraft. *IEEE Transactions on Aerospace and Electronic Systems*, 54(6):2992–3003.

[144] Quero Martin, D., Rabádan, G. J., and Krüger, W. R. (2015). On an Innovative Approach to account for Gust Aerodynamic Nonlinearities. In *AIAA SciTech*, number 5-9 January 2015, Kissimmee, Florida, pages 1–17, Reston, Virginia. American Institute of Aeronautics and Astronautics.

[145] Raghavan, B. and Patil, M. J. (2009). Flight dynamics of high aspect-ratio flying wings: Effect of large trim deformation. *Journal of Aircraft*, 46(5):1808–1812.

[146] Rapinett, A. (2009). Zephyr: A high altitude long endurance unmanned air vehicle.

[147] Rasmussen, C. E. and Williams, C. K. I. (2005). *Gaussian Processes for Machine Learning (Adaptive Computation and Machine Learning)*. The MIT Press.

[148] Reimer, L., Ritter, M., Heinrich, R., and Kr, W. (2015). CFD-based Gust Load Analysis for a Free-flying Flexible Passenger Aircraft in Comparison to a DLM-based Approach. In *22nd AIAA Computational Fluid Dynamics Conference*, number June, pages 1–17.

[149] Reissner, E. (1973). On One-Dimensional Large-Displacement Finite-Strain Beam Theory. *Studies in Applied Mathematics*, 52(2):87–95.

[150] Rodden, W. P., Taylor, P. F., and McIntosh, S. C. (1998). Further Refinement of the Subsonic Doublet-Lattice Method.

[151] Roger, K. L. (1977). Airplane math modelling and active aeroelastic control design. *AGARD Structures and Materials Panel*.

[152] Sarpkaya, T. (1989). Computational methods with vortices—the 1988 freeman scholar lecture. *Journal of Fluids Engineering*, 111(1):5–52.

[153] Schmidt, D. K. (2012). *Modern Flight Dynamics*. Mc-Graw Hill, New York, NY, USA.

[154] Schoor, M. C. V. and Flotow, A. H. V. (1990). Aeroelastic Characteristics of a Highly Flexible Aircraft. *Journal of Aircraft*, 27(10):901–908.

[155] Schuster, D. M., Beran, P. S., and Huttsell, L. J. (1997). Application of the ENS3DAE Euler/Navier-Stokes Aeroelastic Method. Technical report.

[156] Sebastian, T. and Lackner, M. (2012). Development of a free vortex wake method code for offshore floating wind turbines. *Renewable Energy*, 46:269–275.

[157] Shabana, A. (2020). *Dynamics of multibody systems*. Cambridge university press.

[158] Shabana, A. A. (1985). Substructure synthesis methods for dynamic analysis of multi-body systems. *Computers & Structures*, 20(4):737–744.

[159] Shabana, A. A. (1997). Flexible Multibody Dynamics: Review of Past and Recent Developments. *Multibody System Dynamics*, 1(2):189–222.

[160] Shearer, C. and Cesnik, C. (2006). Modified generalized alpha method for integrating governing equations of very flexible aircraft. In *47th AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference 14th AIAA/ASME/AHS Adaptive Structures Conference 7th*. American Institute of Aeronautics and Astronautics.

[161] Shearer, C. M. and Cesnik, C. E. (2007). Nonlinear Flight Dynamics of Very Flexible Aircraft. *Journal of Aircraft*, 44(5):1528–1545.

[162] Silva, W. (2005). Identification of nonlinear aeroelastic systems based on the volterra theory: progress and opportunities. *Nonlinear Dynamics*, 39(1):25–62.

[163] Silvestre, F. J., Guimaraes Neto A. B, Mendes Bertolin, R., da Silva, R. G., and Paglione, P. (2017). Aircraft Control Based on Flexible Aircraft Dynamics. *Journal of Aircraft*, 54(1).

[164] Simeon, B. (2006). On lagrange multipliers in flexible multibody dynamics. *Computer methods in applied mechanics and engineering*, 195(50-51).

[165] Simo, J. (1985). A finite strain beam formulation. The three-dimensional dynamic problem. Part I. *Computer Methods in Applied Mechanics and Engineering*, 49(1):55–70.

[166] Simo, J. and Vu-Quoc, L. (1986). A three-dimensional finite-strain rod model. part II: Computational aspects. *Computer Methods in Applied Mechanics and Engineering*, 58(1):79–116.

[167] Simo, J. C. and Vu-Quoc, L. (1988). On the dynamics in space of rods undergoing large motions – A geometrically exact approach. *Computer Methods in Applied Mechanics and Engineering*, 66(2):125–161.

[168] Simpson, R. J. and Palacios, R. (2013). Numerical aspects of nonlinear flexible aircraft flight dynamics modeling. In *54th AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference*. American Institute of Aeronautics and Astronautics.

[169] Simpson, R. J., Palacios, R., and Murua, J. (2013a). Induced-Drag Calculations in the Unsteady Vortex Lattice Method. *AIAA Journal*, 51(7):1775–1779.

[170] Simpson, R. J. S. (2015). *Unsteady aerodynamics, reduced-order modelling, and predictive control in linear and nonlinear aeroelasticity with arbitrary kinematics*. PhD thesis, Imperial College London.

[171] Simpson, R. J. S., Palacios, R., and Murua, J. (2013b). Induced drag calculations in the unsteady vortex lattice method. *AIAA Journal*, 51(7).

[172] Smagorinsky, J. (1963). General circulation experiments with the primitive equations. *Monthly Weather Review*, 91(3):99–164.

[173] Smith, M., Patil, M., and Hodges, D. (2001a). Cfd-based analysis of nonlinear aeroelastic behavior of high-aspect ratio wings. In *19th AIAA Applied Aerodynamics Conference*, Ancheim, California.

[174] Smith, M. J. (1995). Computational considerations of an Euler/Navier-Stokes aeroelastic method for a hovering rotor. In *33rd Aerospace Sciences Meeting and Exhibit*.

[175] Smith, M. J., Patil, M. J., and Hodges, D. H. (2001b). CFD-based analysis of nonlinear aeroelastic behavior of high-aspect ratio wings. In *42nd AIAA Structures, Structural Dynamics, and Materials Conference and Exhibit*.

[176] Smith, M. J., Schuster, D. M., Huttsell, L., and Buxton, B. (1996). Development of an Euler/Navier-Stokes Aeroelastic Method for Three-dimensional Vehicles with Multiple Flexible Surfaces. In *AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics and Materials Conference*, pages 1733–1746.

[177] Snoek, J., Larochelle, H., and Adams, R. P. (2012). Practical bayesian optimization of machine learning algorithms. In *Advances in neural information processing systems*.

[178] Stanford, B. and Beran, P. (2013). Direct flutter and limit cycle computations of highly flexible wings for efficient analysis and optimization. *Journal of Fluids and Structures*, 36:111 – 123.

[179] Stanford, B. K. and Beran, P. S. (2010). Analytical sensitivity analysis of an unsteady vortex-lattice method for flapping-wing optimization. *Journal of Aircraft*, 47(2):647–662.

[180] Su, W. and Cesnik, C. E. S. (2011). Dynamic response of highly flexible flying wings. *AIAA Journal*, 49(2):324–339.

[181] Suciu, E. O. and Morino, L. (1977). Nonlinear steady incompressible lifting-surface analysis with wake roll-up. *AIAA Journal*, 15(1):15–17.

[182] Tang, D. and Dowell, E. (2004). Effects of geometric structural nonlinearity on flutter and limit cycle oscillations of high-aspect-ratio wings. *Journal of Fluids and Structures*, 19(3):291 – 306.

[183] Tang, D. and Dowell, E. H. (2001). Experimental and theoretical study on aeroelastic response of high-aspect-ratio wings. *AIAA Journal*, 39(8).

[184] Taylor, G. I. (1938). The spectrum of turbulence. *Proceedings of the Royal Society of London. Series A-Mathematical and Physical Sciences*, 164(919):476–490.

[185] Teixeira, P. C. and Cesnik, C. E. S. (2017). Inclusion of Propeller Effects on Aeroelastic Behavior of Very Flexible Aircraft. In *International Forum on Aeroelasticity and Structural Dynamics (IFASD)*, number 6.

[186] The GPyOpt authors (2016). Gpyopt: A bayesian optimization framework in python. http://github.com/SheffieldML/GPyOpt.

[187] Thormann, R. and Widhalm, M. (2013). Linear-Frequency-Domain Predictions of Dynamic-Response Data for Viscous Transonic Flows. *AIAA Journal*, 51(11):2540–2557.

[188] Tijdeman, H. (1977). Investigations of the transonic flow around oscillating airfoils.

[189] Tozer, T. C. and Grace, D. (2001). High-altitude platforms for wireless communications. *Electronics Communication Engineering Journal*, 13(3):127–137.

[190] Valavanis, K. P. (2008). *Advances in unmanned aerial vehicles: state of the art and the road to autonomy*, volume 33. Springer Science & Business Media.

[191] Van Schoor, M. C. and Von Flotow, A. H. (1990). Aeroelastic characteristics of a highly flexible aircraft. *Journal of Aircraft*, 27(10):901–908.

[192] Vassberg, J. and Vassberg, J. (1997). A fast surface-panel method capable of solving million-element problems. In *35th Aerospace Sciences Meeting and Exhibit*. American Institute of Aeronautics and Astronautics.

[193] Vazquez, E. and Bect, J. (2010). Convergence properties of the expected improvement algorithm with fixed mean and covariance functions. *Journal of Statistical Planning and inference*, 140(11):3088–3095.

[194] Veers, P. S. (1988). Three-dimensional wind simulation. Technical report, Sandia National Laboratories. SAND88-0152 UC-261.

[195] Veubeke, B. D. (1976). The dynamics of flexible bodies. *International Journal of Engineering Science*, 14(10):895–913.

[196] von Kármán, T. (1963). General aerodynamic theory. perfect fluids. *Aerodynamic theory*, 2:328–336.

[197] Voutsinas, S. G. (2006). Vortex methods in aeronautics: how to make things work. *International Journal of Computational Fluid Dynamics*, 20(1):3–18.

[198] Wales, D. J. and Doye, J. P. K. (1997). Global optimization by basin-hopping and the lowest energy structures of lennard-jones clusters containing up to 110 atoms. *The Journal of Physical Chemistry A*, 101(28).

[199] Wang, S.-T. and Frost, W. (1980). Atmospheric turbulence simulation techniques with application to flight analysis.

[200] Wang, Y., Wynn, A., and Palacios, R. (2018). Nonlinear Aeroelastic Control of Very Flexible Aircraft Using Model Updating. *Journal of Aircraft*, 55(4).

[201] Wang, Z., Chen, P., Liu, D., Mook, D., and Patil, M. (2006). Time domain nonlinear aeroelastic analysis for hale wings. In *47th AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference 14th AIAA/ASME/AHS Adaptive Structures Conference 7th*, Newport, Rhode Island.

[202] Wang, Z., Chen, P. C., Liu, D. D., and Mook, D. T. (2010a). Nonlinear-aerodynamics/nonlinear-structure interaction methodology for a high-altitude long-endurance wing. *Journal of aircraft*, 47(2).

[203] Wang, Z., Chen, P. C., Liu, D. D., and Mook, D. T. (2010b). Nonlinear-Aerodynamics/Nonlinear-Structure Interaction Methodology for a High-Altitude Long-Endurance Wing. *Journal of Aircraft*, 47(2):556–566.

[204] Wasfy, T. M. and Noor, A. K. (1996). Modeling and sensitivity analysis of multibody systems using new solid, shell and beam elements. *Computer Methods in Applied Mechanics and Engineering*, 138(1-4):187–211.

[205] Wasfy, T. M. and Noor, A. K. (2003). Computational strategies for flexible multibody systems. *Applied Mechanics Reviews*, 56(6):553.

[206] Waszak, M. R. and Schmidt, D. K. (1988). Flight dynamics of aeroelastic vehicles. *Journal of Aircraft*, 25(6):563–571.

[207] Wright, J. R. and Cooper, J. E. (2008). *Introduction to aircraft aeroelasticity and loads*, volume 20. John Wiley & Sons.

[208] Wyngaard, J. C. (2010). *Turbulence in the Atmosphere*. Cambridge University Press.

[209] Xie, C., Liu, Y., Yang, C., and Cooper, J. (2016). Geometrically nonlinear aeroelastic stability analysis and wind tunnel test validation of a very flexible wing. *Shock and Vibration*, 2016.

[210] Xie, S., Ghaisas, N., and Archer, C. L. (2015). Sensitivity issues in finite-difference large-eddy simulations of the atmospheric boundary layer with dynamic subgrid-scale models. *Boundary-Layer Meteorology*, 157(3):421–445.

[211] Yagil, L., Raveh, D. E., and Idan, M. (2018). Deformation Control of Highly Flexible Aircraft in Trimmed Flight and Gust Encounter. *Journal of Aircraft*, 55(2):829–840.

[212] Yang, Z. and Mohammed, A. (2010). Wireless communications from high altitude platforms: Applications, deployment and development. In *2010 IEEE 12th International Conference on Communication Technology*, pages 1476–1479.

[213] Zbrozek, J. K. (1965). Atmospheric gusts: Present state of the art and further research. *The Journal of the Royal Aeronautical Society*, 69(649):27–45.

[214] Zhao, Z. and Ren, G. (2011). Multibody Dynamic Approach of Flight Dynamics and Nonlinear Aeroelasticity of Flexible Aircraft. *AIAA Journal*, 49(1):41–54.