PARAMETRIC PARAGLIDER MODELING

A Thesis

presented to

the Faculty of California Polytechnic State University,

San Luis Obispo

In Partial Fulfillment

of the Requirements for the Degree

Master of Science in Electrical Engineering

by

Peter Frank Heatwole

March 2022

COMMITTEE MEMBERSHIP

TITLE: Parametric Paraglider Modeling

AUTHOR: Peter Frank Heatwole

DATE SUBMITTED: March 2022

COMMITTEE CHAIR: Bridget Benson, Ph.D.
Professor of Electrical Engineering

COMMITTEE MEMBER: Lynne Slivovsky, Ph.D.
Professor of Electrical Engineering

COMMITTEE MEMBER: John Oliver, Ph.D.
Professor of Electrical Engineering

ABSTRACT

Parametric Paraglider Modeling

Peter Frank Heatwole

Dynamic simulations are invaluable for studying system behavior, developing control models, and running statistical analyses. For example, paraglider flight simulations could be used to analyze how a wing behaves when it encounters wind shear, or to reconstruct the wind field that was present during a flight. Unfortunately, creating dynamics models for commercial paraglider wings is difficult: not only are detailed specifications unavailable, but even if they were, a detailed model would be laborious to create. To address that difficulty, this project develops a paraglider flight dynamics model that uses parametric components to model commercial paraglider wings given only limited technical specifications and knowledge of typical wing design. To validate the model design and implementation, an aerodynamic simulation of a reference paraglider canopy is compared to wind tunnel measurements, and a dynamic simulation of a commercial paraglider system is compared to basic flight test data. The entirety of the models and example wings are available as an open source library [1] built on the Python scientific computing stack.

ACKNOWLEDGMENTS

First and foremost I am grateful to my family and friends for the their unending support during this unexpected journey. You never gave up, and held your tongue when the best laid plans went awry.

To Bridget: my friend and advisor, I couldn't have done it without your encouragement. It's been quite the experience.

Also, as computer engineering student with no background in aerodynamics, it was a long road of self-study to reach even my humble level of understanding of this topic. I would like to extend a warm "Thank you!" to Hervé Belloc for his generosity to share the wind tunnel data, and to Robert Kulhánek for taking the time to explain some subtleties of aerodynamics.

And my enduring respect and gratitude to all the software developers that contributed to the free software ecosystem used to produce this project; this project would be impossible without their work. In particular, I am ever grateful for the scientific computing stack (NumPy, SciPy, matplotlib) and the open source aerodynamics models (XFOIL, AVL, and XFLR5).

CONTENTS

LIST OF TABLES

LIST OF FIGURES

Chapter 1

INTRODUCTION

1.1 Overview

The objective of this paper is to create a set of parametric models that can estimate the flight dynamics of commercial paraglider wings using only limited technical specifications.

In this paper, *modeling* refers to creating a mathematical representation of a physical characteristic or behavior. A *dynamics model* is a mathematical function that computes the acceleration of an object given the forces that act on it, as described by Newton's 2nd law of motion (1.1):

$$\begin{aligned} \text{Translational} \qquad & F = ma \\ \text{Angular} \qquad & M = J\alpha \end{aligned} \tag{1.1}$$

These equations show that to compute the translational acceleration a and the rotational acceleration α, a dynamics model requires:

1. The mass m and mass moment of inertia J
2. The forces F and moments M

For a paraglider, the forces and moments that act on it are determined by its current velocity, the relative wind flowing past the glider, air density, gravity, and the pilot control inputs. The motion that is produced are the *flight dynamics*, and the equations that represent how those inputs produce the accelerations are called a *flight dynamics model*:

Fig. 1.1: Flight dynamics model block diagram

The purpose of these *flight dynamics models* is to enable *dynamic simulations*. A *dynamic simulation* is when acceleration is integrated over time to produce a record of the

object's velocity and position. The ability to simulate a system's behavior provides opportunities such as studying that behavior, developing control models, and running statistical filtering pipelines. In fact, the inspiration for this project was a question whether statistical flight reconstruction could be used to recreate the wind fields present during a paraglider flight given only a record of its position, in much the same way as researchers attempted to locate the lost Malaysia Airlines Flight 370 [2].

The steps to producing a dynamic simulation can be summarized as follows:

1. Understand the physical system
2. Model its inertial properties and forces
3. Develop the equations of motion (Newton's 2nd law)
4. Integrate the equations of motion over time

The majority of the work for this project is in step 2 (estimating the inertial properties and forces) because the estimation process requires accurate models of the mass distribution and aerodynamics of each component of the glider.

1.2 Modeling challenges

The existence of this project suggests that existing (and freely available) tools for aircraft simulations are inadequate for simulating paragliders. The reason is that paragliders have a variety of unique characteristics that make them difficult to model using tools built for conventional aircraft:

1. Highly curved shape

   Aerodynamics models must simplify the Navier-Stokes equations in order to produce a tractable system of equations. Those simplifications frequently make them incapable of representing the flow field around a nonlinear wing.

2. Low airspeed

Paraglider airspeeds are typically in the range 24–72 [km/h]. They also have relatively short wing sections, with chord lengths ranging from 0.5–3 [m]. These characteristics combined with the reduced airspeed at the inside wingtip during a turn means that the canopy (and the wing tips in particular) are frequently operating at Reynolds values in the 300k range, far below the $Re = 10^6$ range where where viscous effects start to become significant.

3. High angles of attack

   Compounding the issue of operating at low Reynolds values, paragliders frequently operate at high angles of attack, leading to flow separation and the dramatic nonlinear aerodynamic behavior that results. As they approach stall conditions, simple aircraft simulators that rely on linear aerodynamics can dramatically overestimate the true lift produced by the wing.

4. Flexible

   Paragliders are constructed from flexible nylon sheets and rely on air pressure and suspension lines to maintain their shape. Their internal cells billow and wrinkle while the canopy twists and bends in the wind. It can even collapse entirely. Systems that rely on a predetermined geometry are fundamentally incapable of modeling such behavior.

5. Air intakes

   To produce the internal pressure that forms the canopy, paragliders use air intakes at the leading edge which pressurize its volume. These air intakes violate the expected pressure gradients predicted by analyses that use the idealized airfoils used to define the section profiles. As a result, theoretical aerodynamic coefficients underestimate the section drag.

6. Lightweight

   A paraglider canopy is a large volume with a small amount of solid mass. Its low density means that a naive application of Newton's 2nd law will overestimate acceleration

because it fails to account for the momentum of the fluid surrounding the glider, an effect known as *apparent mass*.

In addition to these characteristics, there is another issue that is relatively unique to gliding aircraft:

7. Pilots care about the details of the wing behavior in non-uniform wind fields.

   The reason is that glider pilots rely on the ability to determine the structure of the wind field by sensing the imbalanced forces produced by differences in relative wind vectors across the wing.

Each of these characteristics introduce modeling challenges. The modeling requirements will depend on which of these characteristics the dynamics model attempts to capture.

1.3 Modeling requirements

The nuances of paraglider behavior are dominated by subtle interactions. The design philosophy for this project was to avoid simplifying assumptions whenever reasonable to avoid accidentally masking those subtle interactions. This approach was driven by a desire to answer questions such as:

- How much drag comes from each individual component?
- How important are section-specific Reynolds values?
- How important is apparent mass?
- How does a paraglider react when one side of the wing is in a stronger thermal than the other side?

The desire for accuracy must be balanced with practical limitations, choosing which characteristics to include and which to simplify away. Having considered the tradeoffs, this

project chose the following set of modeling requirements, beginning with the fundamental challenges of the previous section:

1. The aerodynamics method must use the true, nonlinear geometry. It must not flatten the canopy geometry in any dimension.

2. The aerodynamics method must support variable Reynolds values.

3. The aerodynamics method must provide graceful degradation as it approaches high angles of attack. (A decrease in accuracy is acceptable, but assuming linear aerodynamics up to high alpha is not. The goal is to fly the wing into strong thermals which will rapidly increase angle of attack, so the method must at least approximate those conditions.)

4. Canopy deformations due to flexibility will be neglected. This means that glider controls that use non-brake-line manipulations will also be neglected (since they rely on canopy deformations).

5. The aerodynamics method must support empirical viscous correction factors to mitigate the issues caused by a mismatch between the theoretical and actual section profiles.

6. The system model must support apparent mass (in order to verify its significance).

7. The aerodynamics method must support non-uniform vectors along the span.

In addition to those characteristic behaviors, this project had an additional goal:

8. Computationally fast

   The fundamental goal of this project is to enable people to create models of commercial paraglider wings, and that process requires iteration, so the software should pursue simulation speed that would allow rapid iteration.

1.4 Roadmap

The majority of this work is spent producing the models that estimate the inertial properties and resultant forces for each component, but it also develops the additional models necessary to generate flight simulations. For reference, a complete flight simulation architecture is shown in Fig. 1.2. This paper will develop everything inside the "State dynamics" block.

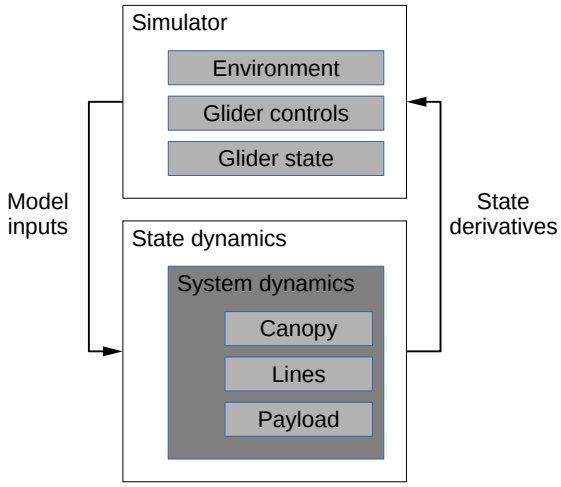

Fig. 1.2: Flight simulation block diagram

The modeling process begins by developing a novel *Foil geometry* with increased flexibility compared to other open source wing modeling tools, enabling simple, parametric representations of typical paraglider canopies. It then chooses a *Foil aerodynamics* method that satisfies those *Modeling requirements* that relate to the canopy aerodynamics. Next, it develops a set of parametric *Component models* using parametrizations that simplify creating models of commercial paraglider systems. Finally, *System dynamics* models combine the components into complete flight dynamics models, and *State dynamics* shows how to define the derivatives of a set of state variables in terms of those system dynamics. Having completed the model derivations, the paper provides a complete *demonstration* of how they can be used to model a commercial paraglider wing. The penultimate chapter provides *Validation* data of the aerodynamics method by comparing wind tunnel measurements for a scale-model paraglider wing against simulated results, as well as comparing simulated polar

curves for the *demonstration* model against basic flight test data. Finally, the *Conclusion* revisits the questions from the *Modeling requirements* and proposes how this material may be used in future work.

Chapter 2

RELATED WORKS

2.1 Flight simulation

This paper develops paraglider flight dynamics models that can be used for flight simulation, which means that this paper is built on the foundations of flight simulation. Flight simulation is simply the specific name of a dynamic simulation that involves a flight dynamics model, and developing a flight dynamics model follows the structure outlined in the *Overview*: understand the system, model the inertia and forces, develop the equations of motion, and integrate them over time.

The first step to creating a model of an aircraft is a familiarity with the physical system and how it behaves. Key concepts in the context of this paper include characteristics of wing geometry; conventions for axes and relative motion; flow angles (angle of attack and sideslip); aerodynamic coefficients; and control inputs, actuators, and surfaces. An approachable starting point is [3], which provides a thorough discussion of the terminology and significance of the major wing design characteristics. Another ubiquitous resource is [4], which may be more suitable to in-depth study.

Next, to model a behavior you must be able to explain the behavior. The unique characteristic of aircraft dynamics is that they experience aerodynamic forces due to their motion relative to the air. The aerodynamic forces on the surfaces of an aircraft are the results of the geometry, relative motion, and characteristics of the fluid. Key concepts include the characteristics of the flow (inviscid versus viscous, laminar versus turbulent, compressibility, etc) and the modeling intuition of Prandtl's seminal work on boundary layers [5] (both 2D and 3D, which are vital to understanding some of the aerodynamic difficulties in simulating flow around a paraglider canopy). When selecting and working with aerodynamics models, it is highly beneficial to have a general awareness of the complexity of Navier-Stokes, and how the variety of aerodynamics models are the result of attempts to produce tractable systems of equations by applying different simplifying assumptions. An excellent introduction to these topics is [6], which provides an approachable introduction to

the underlying physics, overviews of the core aerodynamic models, and how they're derived. Another prevalent work is [7] (or any of Anderson's works). For more targeted discussions, [8] provides clear insight into the theoretical details of common aerodynamics models, and [9] provides a guide to their computational aspects. For a less conventional approach, [10] provides a unique perspective of these aerodynamics models and the assumptions that underlie them, including an excellent discussion of some issues with the NLLT that may shed light on the difficulties that arise when using that approach.

Once the inertial properties, forces, and moments can be determined, they must be synthesized into a complete system dynamics model, which in this case are known as the *equations of motion*. Unlike the simple equations in the *Overview*, the equations describing the translational and angular accelerations of an aircraft cannot always be decoupled; the equations must be solved simultaneously. Producing the equations of motion when such relationships exist involves writing equations for the translational and angular momentum of the system and taking their derivatives with respect to time (since acceleration is the time rate of change of momentum). For a thorough explanation with a focus on aircraft dynamics see [11]; although the notation can be opaque, it provides an excellent development for conservation of momentum of multi-body systems, which is especially useful for understanding the derivations of system models that include degrees of freedom between the paraglider harness and the rest of the system.

Once the equations of motion are known, they can be used to generate simulated trajectories of the aircraft in response to different environmental and pilot inputs. Key concepts include the choice of state variables, coordinate systems and their relative advantages, encoding geometric orientation, representing the environment, and applying numerical integration to the equations of motion to produce the simulated result. For this work I found a complete reference in [12]; the opening chapters provide a masterful introduction to these key concepts, including a principled mathematical notation (adopted by this paper, see *Notation and Symbols*) and a thorough review of vector calculus (especially the counter-intuitive results of taking the derivative of a vector with respect to an accelerating reference frame, which is important when defining the *State dynamics*).

2.2 Paraglider modeling

In addition to the general knowledge of aircraft behavior, it is necessary to understand the unique characteristics of paraglider flight. For practical knowledge, recreational pilot materials make excellent resources. One thorough introduction targeting beginner pilots [13] provides a tour of the components of a paraglider, their function, behavior, and an admirable review of their aerodynamics; if any of the paraglider-specific terminology in this paper is unclear, this book will likely clear up the confusion.

Beyond recreational sources, academic literature relevant to paraglider modeling is typically from one of two branches: parafoil-payload systems, and paragliders. Parafoil-payload systems usually (but not always) refer to large-scale ram-air gliding parachutes intended for heavy payload applications such as cargo delivery and vehicle-recovery (such as landing the X-38 experimental space plane [14], or the more recent work by SpaceX to catch rocket fairings on a boat), while the term "paraglider" usually (but not always) refers to the recreational aircraft. Although the physical characteristics of parafoil-payload systems differ significantly from paragliders due to their scale, carrying capacity, and control schemes, their similarities make much of the research informative, albeit not directly applicable. As a result this section will mix the two groups, noting their differences when significant. Also, as this project has chosen to neglect the effects of canopy deformations, research into modeling those deformations will not be discussed.

The first topic of research is on the aerodynamics of arched, inflatable wings. Their nonlinear geometry made analyses difficult, so early studies were limited to their longitudinal dynamics (fore-aft two-dimensional motion). Alternatively, simple models of their 3D dynamics divide the wing into several discrete segments that act independently (thus neglecting the 3D flow interactions of a real 3D model) [15]. Attempts to account for the full 3D aerodynamics typically involved either measuring the longitudinal and lateral aerodynamic coefficients experimentally [16], or estimating them using vortex lattice and panel methods that can account for their nonlinear geometry by neglecting viscous effects. The significance of the viscous effects led to attempts to incorporate experimental aerodynamic

coefficients via extended lifting-line models; two important works regarding this approach were [17] and [18], which could estimate the 3D aerodynamics of wings with circular arcs, but were unable to account for sweep. As nonlinear lifting-line theory (NLLT) models continue to be developed, their applicability to paraglider wings has greatly improved [19]; for example, [20] successfully applied the method from [21] to a reference paraglider wing in a static flight test, confirming the merit of the of a modern NLLT to this application.

Another significant characteristic of paraglider canopies is their low density, which makes them sensitive to the effects of *apparent mass* [22]. Early attempts to model the apparent mass of a paraglider simplified the wing as an ellipsoid with a single center of rotation [23]. Further developments recognized the inadequacies the ellipsoid model, and adjusted the estimates to account for two separate centers of rotation for rolling and pitching motions [24]. Both models are limited by their assumption of steady flow [25] so their adequacy for simulations involving dynamic maneuvers is unclear; nevertheless, the adapted model is assumed to be adequate for the purposes of this paper.

The last major topic of research is the system model. There are many system models in literature, but their key differentiating factors in the context of this project are whether they incorporate apparent mass and how they model the attachment of the harness to the suspension lines. The inclusion of apparent mass appears to be a modeling decision driven by whether the author expected the effect to be significant; papers that exclude apparent mass do so without explicit justification. For the harness connection, models are categorized by their *degrees of freedom* (DoF) and the character of the connection points; a 6-DoF model does not allow the payload to move at all, a 7-DoF allows the payload to translate or rotate (relative to the suspension lines) in one dimension, an 8-DoF adds two degrees of freedom, etc. For a general understanding of the impact, [26] provides a comparative analysis of a fixed (6-DoF) model versus a 9-DoF system model. For a more thorough review of the many available system models, [27] has a seemingly exhaustive list of the models through 2005, including a discussion of those models that account for apparent mass. Two informative models that incorporate apparent mass are [15] (which used the older method in [23]) and [28] (which used the adapted apparent mass model from [24]).

In addition to topical works, there have been several more comprehensive studies. The best place to start is [29]: although it has a parafoil-payload perspective, this approachable paper is a thorough introduction to the terminology, geometric parameters, choice of airfoil, and control schemes of parafoils (which it calls a "ram-air parachute"); this paper also used geometric simplifications to study the canopy aerodynamics and drag contributions, and developed linear models of the longitudinal and lateral dynamics to study performance and stability. Next, for a paraglider perspective, [30] provides a compact survey on the sources of aerodynamic drag; it reviews the impacts of arc, flexibility, air intakes, lines, and pilot. Worth reading immediately after is [20], as it is essentially an updated revision of [30].

The most comprehensive work on paraglider flight dynamics to date is the dissertation [31] that inspired the general structure of this paper. First, it provides an overview of paraglider geometry, construction, and behavior. It then develops a foil geometry that uses the locus of quarter-chord points to position the sections, as well as intuitive parametric definitions of the underlying paraglider canopy structure. For the paraglider components, it develops a model to position the harness as a function of the accelerator control, a continuous brake deflection distribution using both brakes, and the spherical harness model used by this paper. Next, for the canopy aerodynamics it develops a pseudo-LLT (which it acknowledges is an approximation in deference to the project's primary focus on stability and control) using constant 2D aerodynamic coefficients. From the complete aerodynamics model, it then estimates the 3D aerodynamic coefficients and stability derivatives for a linearized model that is used for the remainder of the work, which is focused on performance aspects (such as glide ratio versus equilibrium pitch angle), stability analyses (such as longitudinal stability versus riser position, and roll stability versus sideslip), and controllability (takeoff, maneuvering, and landing).

2.3 This work

As mentioned in the previous paragraph, this project began with [31] as its starting point. While attempting to use those models to recreate commercial paraglider wings, this work identified a collection of improvements that led to newly derived models.

First, it improves the canopy geometry by developing a novel foil geometry model inspired by a suggestion in [32] that allows independent reference points for the x- and yz-positions. This increased flexibility allows accurate representations of existing wings using simple parametric equations, which this work uses to replace the parametric design curves in [31] with new parametrizations that are easier to estimate for an existing paraglider canopy. It also replaces the approximate inertia calculations for the canopy surface and volume with a mesh-based method that can account for different upper and lower surface densities, and the extra solid mass from vertical ribs.

For the canopy aerodynamics, it replaces his pseudo-LLT with a full NLLT ([21], [33]) that supports arbitrary arc, sweep, twist, specific (nonlinear model) aerodynamic coefficients for each section as a function of Reynolds number and deflection distance, and non-uniform wind vectors along the span. Also, instead of modeling trailing edge deflections as section rotations (by adding the deflection angle to the section angle of attack, effectively shifting the coefficient curves), this model uses section coefficients generated from the actual deflected geometry, and accounts for the effects of Reynolds number.

Next, it completely redesigns the suspension line model, keeping only the intuition to replace the "rigging angle" with a displacement vector in the body axes. The new model improves the representation of the brakes by first calculating the deflection distance before calculating the true change in angle of attack (which depends on the section chord), as well as improving the accuracy of the deflection distribution itself. The new model improves the representation of the accelerator by parametrizing the fore and aft connection points instead of fixing them at the leading and trailing edge of the canopy, thus allowing accurate models of commercial wings. Lastly, the new model moves the line drag away from canopy centroid and distributes it into lumped points that can model asymmetric forces between

each semispan.

For the harness, the only minor change was to separate the weight shift distance from an absolute distance to a proportional one controlled by a harness parameter for the maximum displacement. Although functionality equivalent, I personally felt that this change makes simulation scenarios easier to write and understand.

For the system model, this paper derived 6-DoF and 9-DoF models (the 9-DoF is a rederivation of the model used in [34] and [35]) that may optionally incorporate the apparent mass estimates from [24]. The 9-DoF model is included for demonstration and testing purposes, and is not used in any analyses.

The implementation of all models are available as an open source library [1], including example wing models, and the simulations used in this paper are available as part of the open source materials used to produce this paper.

Chapter 3

FOIL GEOMETRY

The essential components of any flying object are the lifting surfaces, or *foils*: by redirecting airflow, a foil exchanges momentum with the air, producing a lifting force that allows the object to fly. The dynamics of a foil depend on its inertial properties and its aerodynamics, both of which can be estimated from its shape.

A foil geometry model describes the shape of a foil by defining the positions of all the points on the foil's surfaces. Although those positions can be defined as an explicit set of points (with interpolation in between), it is much more convenient to decompose them into a set of variables that represent distinct characteristics of the foil's shape. Similarly, those variables may be defined using explicit values, but it is much more convenient to define them using *design curves*: parametric functions that encode that underlying structure of the foil with a small number of intuitive parameters.

This decomposition is essential to this project, because the foils of interest are commercial paraglider wings, and manufacturers do not provide explicit geometry data; at best, marketing materials and user manuals provide basic summary specifications, which means the majority of the geometry is unknown. Generating a surface model from summary information requires making educated guesses about the missing structure in order to generate a complete geometry. That assumed structure takes the form of domain expertise encoded in the design curves, which augment the summary data to produce a fully specified model.

The difficulty with this approach is that the choice of variables in a geometry model controls how a designer must specify the structure. More variables increase model flexibility at the cost of increased complexity, so the goal is to choose the smallest number of variables that provide the designer with adequate flexibility. Existing foil models are inflexible, making strong assumptions about how foils are most naturally defined, and that inflexibility forces the remaining complexity into the design curves. This unnecessary complication makes it difficult to describe a parafoil using simple parametric functions: they must not only encode the fundamental structure, they must also translate that structure into the

variables that define the model. Instead of the geometry model adapting to the needs of the design curves, the design curves must adapt to the inflexibility of the model.

The solution developed in this chapter is to reject the assumption that predefined reference points are the most convenient way to position the elements of a foil surface. The result is a novel foil geometry that fully decouples the design curves, allowing each variable to be designed independently. It also presents a simplified model that eliminates most of the additional complexity of the expanded model. The simplified model is both flexible and intuitive for designing highly nonlinear foil geometries (such as paraglider canopies) using simple parametric functions.

But first, a remark on notation: in this chapter, the lifting surface of an aircraft is referred to as a *foil* instead of using the conventional terms *wing* or *canopy* (for traditional aircraft or parafoils, respectively). This unconventional term was chosen to avoid two generalization issues. First, although *wing* is the conventional term for the primary lifting surfaces of non-rotary aircraft, the paragliding community already uses the term *paraglider wing* to reference not only the lifting surface but also the supporting structure connected to it, such as suspension lines, risers, etc. Second, although this project is primarily concerned with parafoils, the content in this chapter is not limited to parafoil canopies, making "canopy" a poor choice.

In addition, note that these are idealized geometry models, not detailed structural models. Structural models include physical details that can be used to simulate effects such as internal forces and wing deformations [36]. Unfortunately, as discussed earlier, such details are not available for commercial paraglider wings, and such analyses would be time prohibitive even if they were. Instead, this design will model only those details of the shape that can be approximated from the available data. It does not model internal structures, in-flight deformations, or surface deviations from the idealized design target.

3.1 Modeling with wing sections

At its most basic, a foil geometry is the surface of a volume. Points on the surface can be defined with explicit coordinates, or they can be generated using functions that encode aspects of the surface's structure. Explicit geometries are extremely flexible (since they can encode arbitrary amounts of detail), but refining an explicit mesh can be very time consuming (in addition to requiring highly detailed geometry data). Conversely, parametric geometries model the surface mesh indirectly using parametric functions which encode structural knowledge of the shape. In effect, the parameters summarize the structure: a structural parameter communicates more information than an explicit coordinate, which means less work (and less data) is required to specify a design.

The standard first step towards parametrizing a foil geometry is to define it in terms of *wing sections* ([37]; [6], Sec. 5.2). The foil is modeled as a sequence of *sections* (typically arranged spanwise, left to right) over some continuous *section index* s. Each section is assigned a 2D cross-sectional profile, called an *airfoil*, which lies perpendicular to the local spanwise axis. Each airfoil is scaled, positioned, and oriented to produce the *section profile*. Together, the section profiles produce a continuous surface that defines the complete 3D volume.

Wing design using airfoils is thus decomposed into two steps:

1. Specify the scale, position, and orientation of each section
2. Specify the airfoil at each section

In some literature [3] these two steps are described as designing the *planform* and the *profile*, but this description is problematic due to inconsistent uses of the term *planform* across literature. Specifically, in some cases the planform is the complete surface produced by the section chords, and in others "planform" refers to a projected-view of the chord surface onto the xy-plane. Due to this ambiguity, this paper avoids the term *planform* in preference of explicit references such as *chord surface*, *mean camber surface*, or *profile surface*.

Fig. 3.1: Wing section profiles.

Note that section profiles are not the same thing as the ribs of a parafoil. Parafoil ribs are the internal structure that produce the desired section profile at specific points along the span.

3.1.1 Section index

In order to generate a foil from discrete wing sections (and to support queries about their individual properties) each section must be assigned a unique identifier which this paper refers to as a *section index* s. This term is deliberately generic. Some aeronautics literature use the term *spanwise station*, but "spanwise" is ambiguous: some papers use "spanwise" to refer to the absolute y-coordinate of some reference point embedded in each section, while others refer to the linear distance along the curve through those reference points. The term *section index* generalizes these concepts and provides an arbitrary reference to any choice of unique identifier over the set of sections.

However, avoiding ambiguity is the not the primary purpose of this generality. The real goal is to avoid unnecessary coupling of the design curves that define the geometry. Instead of committing to a definition immediately, delaying the choice of section index allows a designer the freedom to define the section index in terms of the geometry, or the geometry in terms of the section index, or a even a mixture of the two. This freedom will be used later by the *Simplified model* to enable particularly simple parametric design curves.

3.1.2 Airfoil

The building block of each section is its dimensionless cross-sectional profile, called an *airfoil*. The volume of the wing is generated by the continuum of neighboring airfoils, so the choice of 2D airfoils is vital to designing the flow field characteristics over the 3D wing. The choice involves trade-offs specific to the application (for example, thicker airfoils tend to offer more gentle stall characteristics in exchange for a small increase in drag); as a result, the variety of airfoil designs is very diverse.

Fig. 3.2: Airfoils.

Airfoils are conventionally described using terms that assume the airfoil can be divided into upper and lower surfaces. The upper and lower surfaces are separated by two points defined by a straight *chord line* that runs from the rounded leading edge back to the sharp trailing edge. The curve created by the midpoints between the upper and lower surface curves is the *mean camber line*.

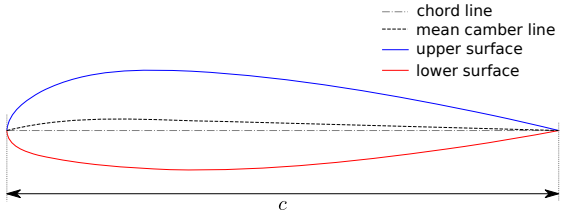

Fig. 3.3: Components of an airfoil.

Another standard design parameter for an airfoil is its *thickness distribution*. Unfortunately, the mean camber line and thickness distribution are not universally defined, because there are two conventions for measuring the airfoil thickness: perpendicular to the chord line (sometimes referred to as the "British" convention), or perpendicular to the mean camber line (the "American" convention). The thickness convention also determines what point is designated the *leading edge*. For the "British" convention the leading edge is the point where the curve is perpendicular to a line from the trailing edge. For the "Ameri-

can" convention, the leading edge is the "leftmost" point with the smallest radius (greatest curvature).

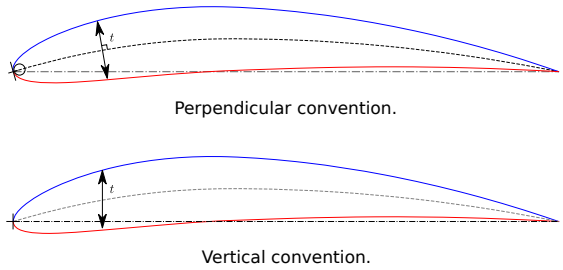

Fig. 3.4: Airfoil thickness conventions.

As a result, the exact value of the mean camber line and thickness depends on the thickness convention, but in general the mean camber line will lie halfway between an upper and lower surface whose separation distance is specified by the thickness distribution. Fortunately, this ambiguity is irrelevant except when comparing airfoil design parameters.

3.1.3 Scale

By convention, airfoils are normalized to a unit chord length. Similarly, the aerodynamic coefficients associated with an airfoil are also dimensionless. To generate the geometry and compute the aerodynamic forces associated with a wing segment, both the airfoil and its aerodynamic coefficients must be scaled in units appropriate to the model.

Although conceptually simple, section scale plays a large role in controlling the aerodynamic behavior of a wing segment; in fact, all but the most basic foils have variable section chord lengths. The only fundamental requirement is that the sections collectively produce enough aerodynamic lift to support the aircraft, but beyond that a foil designer is free to use use spanwise variation to control behavior such as:

- Spanwise loading (the chord lengths are one factor, along with choice of section profile and orientation/twist, that can be used to encourage an elliptical load distribution, thus minimizing induced drag)

- Weight distribution

- Relative importance of wing segments (if the wingtips are smaller then they contribute less to the loading, making the loading is less sensitive to wingtip stalls, leading to "gentler" stall characteristics)

3.1.4 Position

The relative position of the sections is fundamental to controlling important foil characteristics such as *span*, *sweep*, and *arc* [3]. Span (the width of the wing, roughly speaking) together with the chord distribution determines the *aspect ratio* of a foil, which impacts characteristics such as aerodynamic efficiency and maneuverability. Sweep (the fore-aft relative positioning of the sections) is important for controlling spanwise airflow. Arc (the vertical relative positioning of the sections, roughly speaking) is primarily used to increase the roll stability of conventional wings, although for parafoils the *arc anhedral* is essential to designing the spanwise loading across the suspension lines.

To define their layout, each section must be positioned by specifying a vector in foil coordinates of some *reference point* in the section's local coordinate system. For example, the most common choice of reference point is the leading edge of the section profile; by convention the section leading edge will coincide with the origin of the airfoil coordinate system, which means no additional translations are required to position the profile. This conventional but inflexible choice is demonstrated by the *Basic model*, then relaxed by the *Expanded model*, and made convenient by the *Simplified model*.

3.1.5 Orientation

The last degree of freedom for a wing section is its orientation. Instead of pointing straight ahead, the can roll and twist to change their angle of attack in different flight conditions. Changing the wind angles affects both their aerodynamic coefficients as well as the direction of the force and moment produced by that section. Controlling the strength, magnitude, and orientation of the section forces can be used to control characteristics such as the zero-lift angle of the wing, spanwise loading (the lift distribution, which also affects the induced

drag of the wing), stall profile (how stall conditions develop across the span), and dynamic stability (such as the roll-yaw coupling exhibited by wings with arc anhedral).

3.2 Basic model

Choosing to model a foil using *wing sections* means that the surfaces are defined by 2D airfoils. The 2D airfoil curves must be converted into a 3D section-local coordinate system, then scaled, positioned, and oriented relative to the foil coordinate system. This "basic" model describes how that is done by conventional wing modeling tools, which position the sections by their leading edge.

First, let P represent any point in a wing section (such as points on the chord, mean camber line, or profile), and LE be the leading edge of that section. It is conventional to share the origin between the airfoil and section coordinate systems, and specify the section position using the section leading edge, so using the *notation* of this paper, a general equation for the position of that point P with respect to the foil origin O, written in terms of the foil coordinate system f, is:

$$\boldsymbol{r}^f_{P/O} = \boldsymbol{r}^f_{P/LE} + \boldsymbol{r}^f_{LE/O} \tag{3.1}$$

Assuming the foil geometry is symmetric, designate the central section the foil *root*, and let the 3D foil inherit the 3D coordinate system defined by the root section. Points in section (local) coordinate systems s must be rotated into the foil (global) coordinate system f. Given the *direction cosine matrix* $\boldsymbol{C}_{f/s}$ between the section and foil coordinate systems, position vectors in foil coordinates can be written in terms of section coordinates:

$$\boldsymbol{r}^f_{P/LE} = \boldsymbol{C}_{f/s}\boldsymbol{r}^s_{P/LE} \tag{3.2}$$

Because airfoil curves are defined in the 2D airfoil-local coordinate system a, another transformation is required to convert them into the 3D section-local coordinate system s. The convention for airfoil coordinates places the origin at the leading edge, with the x-axis pointing from the leading edge towards the trailing edge, and the y-axis oriented towards the upper surface. This paper uses a front-right-down convention for all 3D coordinate

systems, so the conversion from 2D airfoil coordinates a to 3D section coordinates s can be written as a matrix transformation:

$$\boldsymbol{T}_{s/a} \stackrel{\text{def}}{=} \begin{bmatrix} -1 & 0 \\ 0 & 0 \\ 0 & -1 \end{bmatrix} \tag{3.3}$$

Next, the airfoil must be scaled. By convention, airfoil geometries are normalized to a unit chord, so the section geometry defined by the airfoil must be scaled by the section chord c. Writing the points in terms of relative position vectors defined in the foil coordinate system produces:

$$\boldsymbol{r}^f_{P/LE} = \boldsymbol{C}_{f/s}\boldsymbol{T}_{s/a}\, c\, \boldsymbol{r}^a_{P/LE} \tag{3.4}$$

The complete general equation for arbitrary points P in each section s is then:

$$\boldsymbol{r}^f_{P/O}(s) = \boldsymbol{C}_{f/s}(s)\boldsymbol{T}_{s/a}\, c(s)\, \boldsymbol{r}^a_{P/LE}(s) + \boldsymbol{r}^f_{LE/O}(s) \tag{3.5}$$

In this form it is clear that a complete geometry definition requires four *design curves* that define the variables for every section:

$$\begin{aligned} c(s) &\quad \text{Scale} \\ \boldsymbol{r}^f_{LE/O}(s) &\quad \text{Position} \\ \boldsymbol{C}_{f/s}(s) &\quad \text{Orientation} \\ \boldsymbol{r}^a_{P/LE}(s) &\quad \text{Airfoil} \end{aligned} \tag{3.6}$$

3.3 Expanded model

The basic equation (3.5) is an explicit mathematical equivalent of the approach used by most freely available wing modeling tools. However, although it is technically sufficient to describe arbitrary foils composed of airfoils, its inflexibility can introduce incidental complexity into what should be fundamentally simple design curves.

For example, consider a delta wing with a straight trailing edge:

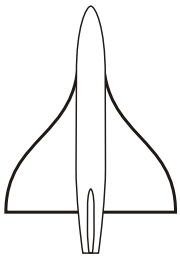

Fig. 3.5: Ogival delta wing planform.

Figure by Wikimedia contributor "Steelpillow", distributed under a CC-BY-SA 3.0 license.

The wing geometry is fundamentally simple. Its specification should be equally simple, but defining this wing with a model that is only capable of positioning sections by their leading edge makes that impossible. Instead, the position curve must be just as complex as the scale function (chord length) in order to achieve the straight trailing edge. The simplicity of the model has forced an artificial coupling between the design curves.

The problem becomes much more severe when section section chords no longer lie in the xy-plane, because the trailing edge position is no longer a simple x-coordinate offset; instead, all of the scale, position, and orientation design curves are coupled together, making design iterations incredibly tedious. Whether the adjustments are performed manually or with the development of additional tooling, the fact is the extra work is unnecessary.

The solution is to decouple all of the design curves by allowing section position to be specified using arbitrary reference points in the section coordinate systems. This can be accomplished by decomposing their positions into two vectors: one from the section *leading edge* LE to some arbitrary *reference point* RP, and one from the reference point to the *foil origin* O:

$$\boldsymbol{r}^f_{LE/O} = \boldsymbol{r}^f_{LE/RP} + \boldsymbol{r}^f_{RP/O} \tag{3.7}$$

Although this decomposition increases model complexity, the additional flexibility allows a designer to choose whichever point in each section's coordinate system will produce the simplest geometry specification. The basic model (3.5) is replaced by an expanded equation

with a new set of design curves:

$$\boldsymbol{r}^f_{P/O}(s) = \boldsymbol{C}_{f/s}(s)\boldsymbol{T}_{s/a}\, c(s)\, \boldsymbol{r}^a_{P/LE}(s) + \boldsymbol{r}^f_{LE/RP}(s) + \boldsymbol{r}^f_{RP/O}(s) \tag{3.8}$$

$c(s)$	Scale
$\boldsymbol{r}^f_{RP/O}(s)$	Position
$\boldsymbol{C}_{f/s}(s)$	Orientation
$\boldsymbol{r}^a_{P/LE}(s)$	Airfoil
$\boldsymbol{r}^f_{LE/RP}(s)$	Reference point

(3.9)

3.4 Simplified model

The *Basic model* is adequate to represent wings arbitrary foils composed of airfoils, but its inflexibility forced incidental complexity into the design curves. The *Expanded model* provides additional flexibility, but it's generality can make it difficult for a designer to identify which aspects of the foil structure result in a simple parametric representation. This section identifies several simplifying assumptions that provide a foundation for a particularly concise representation of many foils (parafoils in particular). The result is an intuitive, partially-parametrized foil geometry model that decouples the design curves and allows a parafoil to be rapidly approximated using only minimal available data, even if that data was obtained from a flattened version of the parafoil.

3.4.1 Section index

Although most tools do not explicitly announce to their choice of section index, there are two conventions in common use: the most common is to use the reference point y-coordinate ($s = y$, or its normalized version $s = \frac{y}{b/2}$). Although simple and intuitive for flat wings, defining a nonlinear geometry in terms of y can become unwieldy, so another common choice is to use the linear distance along the locus of reference points $\boldsymbol{r}_{RP/O}$ (or its normalized version that ranges ± 1). Unfortunately, both are problematic for modeling a paraglider canopy using the most readily-available data.

When trying to create a model of a flexible wing like a paraglider canopy, it is much easier to take measurements when the wing is stretched out flat. When the canopy is flat it is possible to measure $c(s)$ and $x(s)$ directly, whether from the physical wing or from photos (such as are found in user manuals). Also, it is trivial to measure the flattened span compared to trying to measure the span of an in-flight canopy. The solution is to use the normalized section y-coordinates from the flattened foil:

$$s = \frac{y_{flat}}{b_{flat}/2} \tag{3.10}$$

Not only does this choice make the section index easy to measure from a flattened paraglider canopy, but with a careful choice of reference points it also decouples the yz-coordinates of the reference positions ($yz(s)$) from all the other design curves, which is a key aspect of this model's ability to define complex nonlinear foils using simple parametric functions. The next section explains the process in detail, but the key idea (and why this choice of section index is so important) is that using this definition of s and choosing the same chord position for the y and z components of the reference point you can simply "wrap" the flattened paraglider canopy around $yz(s)$ to produce the final geometry. It becomes possible design the flattened foil geometry before designing its arc, a natural process that enables the direct use of the most readily available measurements for commercial paraglider canopies.

3.4.2 Reference point

The *Basic model* positions each section using the section origins (the leading edges). The *Expanded model* allows the sections to be positioned using arbitrary reference points anywhere in the 3D section coordinate systems. Although flexible, the freedom of the expanded model does not address the problem of choosing good reference points.

One intuitive choice is to use points on the section chords, in which case the reference point is a function of a chord ratio $0 \leq r \leq 1$. The chord lies on the negative section x-axis, so a reference point at some fraction r along the chord is given by $\boldsymbol{r}^s_{RP/LE} = -r\,c\,\hat{x}^s_s$ (where $\hat{x}^s_s \stackrel{\text{def}}{=} [1\,0\,0]^T$, the x-axis of section s in that section's local coordinate system). Substituting

$\boldsymbol{r}_{LE/RP} = -\boldsymbol{r}_{RP/LE}$ into (3.8) produces:

$$\boldsymbol{r}^f_{LE/O} = \boldsymbol{C}_{f/s}\, r\, c\, \hat{x}^s_s + \boldsymbol{r}^f_{RP/O}$$

Simple and intuitive, this parametrization captures the choices used by every foil modelling tool reviewed for this project. Models that position sections by their leading edge (XFLR5, AVL, MachUpX) are equivalent to setting $r = 0$. Another (less common [31]) choice is to use the quarter-chord positions, in which case $r = 0.25$. The problem with the constraint that reference points lie on the section chords is that it couples the position functions for all three dimensions. For many foil geometries it can be significantly more convenient to use different chord positions for different dimensions.

For example, suppose an engineer is designing a foil with an elliptical chord distribution and geometric twist, and they wish to place the leading edge in the plane $x = 0$ and the trailing edge in the plane $z = 0$. Although the intuitive specification of this foil would be $x(s) = 0, z(s) = 0$, it cannot be used because it needs to position different points on each section chord: the $x(s) = 0$ design requires $r = 0$, but the $z(s) = 0$ design requires $r = 1$. One of the position curves must be changed, introducing unnecessary complexity to make up for this inflexibility.

For another example, a foil designer may want to arc an elliptical planform such that the y- and z-coordinates of the quarter-chord ($r = 0.25$) follow a circular arc while the x-coordinate of the trailing edge ($r = 1$) is a constant. Because of the elliptical chord distribution, the x-coordinates of the quarter-chord that would produce a straight trailing edge are distinctly non-constant; if geometric twist is present the issue becomes even more severe. What should be a simple $x(s) = 0$ to specify the straight trailing edge must become a complex function with no simple analytical representation.

The underlying problem is that the designer cannot specify their design directly using a shared reference point that lies directly on the chord; instead, they must translate their design into an alternative specification using positions that accommodate the shared reference point.

The solution is that instead of using a shared reference point directly on the chord for all dimensions, allow each dimension to choose independent reference points along the chord, and associate each dimension of the position design curve with that dimension's coordinate of that dimension's reference point. The $x(s)$ design curve specifies the x-coordinate of the reference point for the x-dimension, etc.

Fortunately, providing this flexibility is easier to implement and use than it is to describe. Instead of a shared r for all three dimension, allow an independent r for each dimension of the reference point:

$$\boldsymbol{R} \stackrel{\text{def}}{=} \begin{bmatrix} r_x & 0 & 0 \\ 0 & r_y & 0 \\ 0 & 0 & r_z \end{bmatrix}$$

where $0 \leq r_x, r_y, r_z \leq 1$ are proportions of the chord, as before. The coordinates of the leading edge relative to the reference point are now the displacement of the section origin relative to the $\{x, y, z\}$ components of the $\{r_x, r_y, r_z\}$ positions along the chord. The resulting equation, which allows completely decoupled positioning for each dimension, is surprisingly simple:

$$\boldsymbol{r}^f_{LE/O} = \boldsymbol{R}\boldsymbol{C}_{f/s} c\, \hat{x}^s_s + \boldsymbol{r}^f_{RP/O}$$

This choice of reference point makes the earlier examples trivial to implement. For the first, which was struggling with the fact that geometric twist has coupled the x and z positions is solved with $\{r_x = 0, r_z = 1\}$ (because the foil is flat, every choice of r_y is equivalent). The second example, which was struggling to define an $x(s)$ to achieve a straight trailing edge, the answer is simply $\{r_x = 1, r_y = 0.25, r_z = 0.25\}$. In both cases, the designer is able to specify their target directly, using simple design curves, with no translation necessary. The reason is that (3.10) combined with $r_y = r_z$ means that changing $yz(s)$ does not change the section index; having designed the orientation and fore-aft position $x(s)$ of a section, changing $yz(s)$ will not affect that design. The curves have been decoupled.

$$\boldsymbol{r}^f_{LE/RP} = \boldsymbol{R}\boldsymbol{C}_{f/s} c\, \hat{x}^s_s \tag{3.11}$$

$$\boldsymbol{R} \stackrel{\text{def}}{=} \begin{bmatrix} r_x & 0 & 0 \\ 0 & r_{yz} & 0 \\ 0 & 0 & r_{yz} \end{bmatrix} \tag{3.12}$$

3.4.3 Orientation

The expanded model (3.8) uses a *direction cosine matrix* (DCM) to define the orientation of each section; the problem is how to define that matrix. A natural parametrization of a DCM is a set of three Euler angles $\langle \phi, \theta, \gamma \rangle$, corresponding to roll, pitch, and yaw. The Euler parametrization replaces the $\mathbb{R}^{3\times3}$ matrix with a 3-vector — three parameters — but the structure of typical parafoils can provide further simplifications.

In particular, observe that when a parafoil is flattened out on the ground, the sections are (essentially) vertical, with no relative roll or yaw. Inflating the parafoil and using the suspension lines to form the arc will naturally roll the sections without affecting the section yaw. These observations reveal that the section orientation produced by inflating a parafoil is well approximated by a single degree of freedom, resulting in a minimal parametrization with a single design variable for section pitch $\theta(s)$.

For the section roll $\phi(s)$, observe that inflating the foil to produce the arc does not produce a shearing effect between sections; instead, the sections roll jointly with the arc. This relationship can be encoded using the derivatives of the $\langle y(s), z(s) \rangle$ components of the position curve $\boldsymbol{r}_{RP/O}(s)$:

$$\phi = \arctan\left(\frac{dz}{dy}\right) \tag{3.13}$$

For the section yaw $\gamma(s)$, inflating the parafoil to produce the arc anhedral will roll the sections in the foil's yz-plane and does not affect the section yaw, which remains zero:

$$\gamma = 0 \tag{3.14}$$

The remaining degree of freedom is the rotation about each sections y-axis. This pitch angle $\theta(s)$, conventionally known as *geometric torsion*, is produced when the wing is manufactured, and is not affected when the flattened wing is shaped into its final arched form.

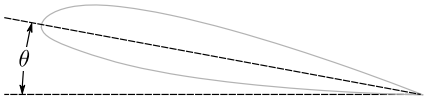

Fig. 3.6: Geometric torsion.

Note that this refers to the angle, and is the same regardless of any particular rotation point.

3.4.4 Summary

In conclusion, the simplifications identified by this model not only reduced the number of parameters of the expanded model (3.9), it also replaced the arbitrary and unwieldy 3D reference points with simple ratios of the section chords. It allows rapid and intuitive conversion of measurements from a flattened paraglider canopy to a foil geometry, and decoupled the design curves to allow the design of each variable to be manipulated without affect the others. In short, it provides the flexibility of the expanded model but without its complexity.

$$\begin{aligned} c(s) &\quad \text{Scale} \\ r_x(s) &\quad \text{Chord ratio for positioning } RP_x \\ r_{yz}(s) &\quad \text{Chord ratio for positioning } RP_y \text{ and } RP_z \\ \boldsymbol{r}^f_{RP/O}(s) &\quad \text{Position} \\ \theta(s) &\quad \text{Pitch} \\ \boldsymbol{r}^a_{P/LE}(s) &\quad \text{Airfoil} \end{aligned} \tag{3.15}$$

3.5 Examples

These examples demonstrate how the simplified model makes it easy to represent nonlinear foil geometries using simple parametric functions, such as constants, absolute functions, ellipticals, and polynomials. For a discussion of the elliptical functions for the arc and chord distributions, see *Parametric design curves*.

All examples show a wireframe view of the chord surface because it is easier to visualize the foil layout. The green dashed lines are projections of the section quarter-chord positions

(shown because of their use in analyzing aerodynamics). The red dashed lines are the projections of the r_x and r_{yz} chord positions.

3.5.1 Delta wing

A delta with with a linear chord distribution and straight trailing edge can be defined with $r_x = 1$ and a piecewise-linear $c(s)$. Unlike conventional wing modeling tools, because the trailing edge is used directly for position in the x-direction, the $x(s)$ curve does not need to be coupled to $c(s)$ to compute offsets for the leading edge.

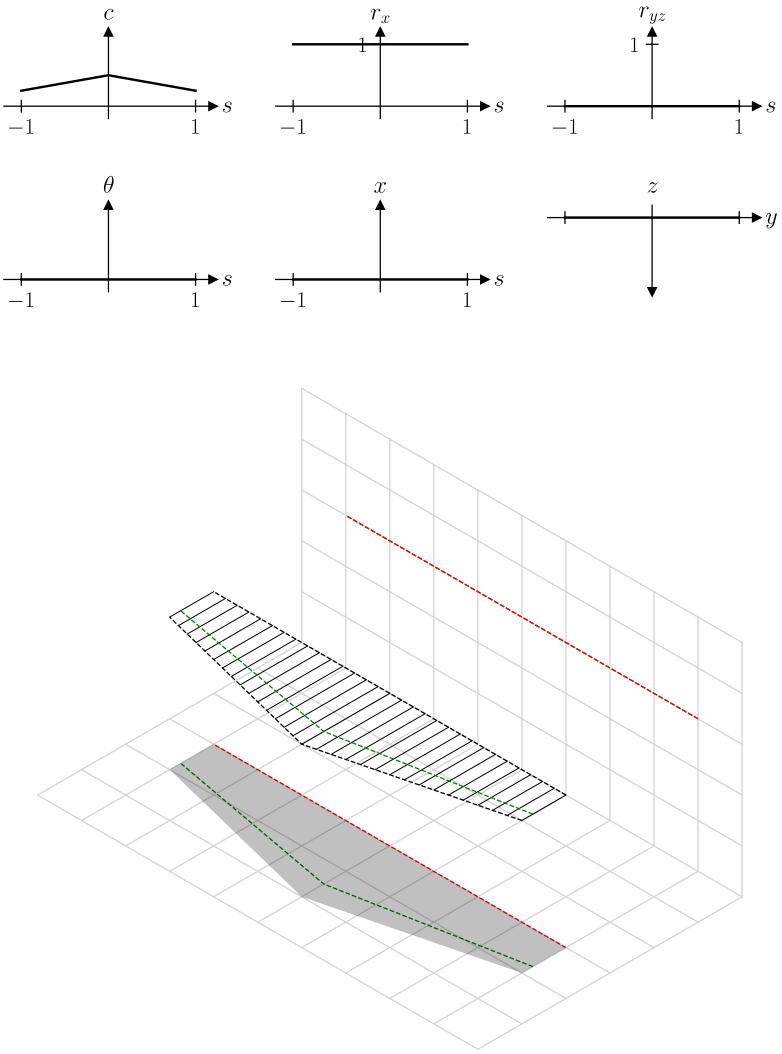

Fig. 3.7: Chord surface of a delta wing planform.

3.5.2 Elliptical wing

Similarly, a flat wing with an elliptical chord distribution and fore-aft symmetric is trivial to define using $r_x = 0.5$ and an elliptical chord function.

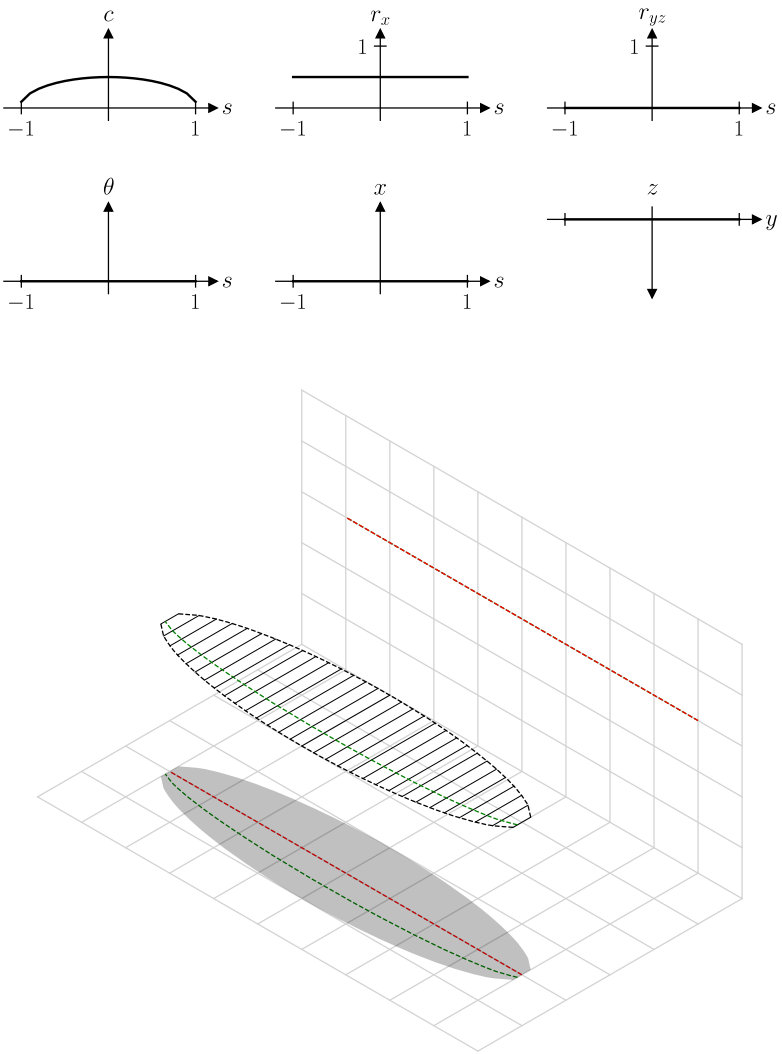

Fig. 3.8: Chord surface of an elliptical wing planform.

3.5.3 Twisted wing

Wings with twist typically use relatively small angles that can be difficult to visualize. Exaggerating the angles with extreme torsion makes it easier to see the relationship.

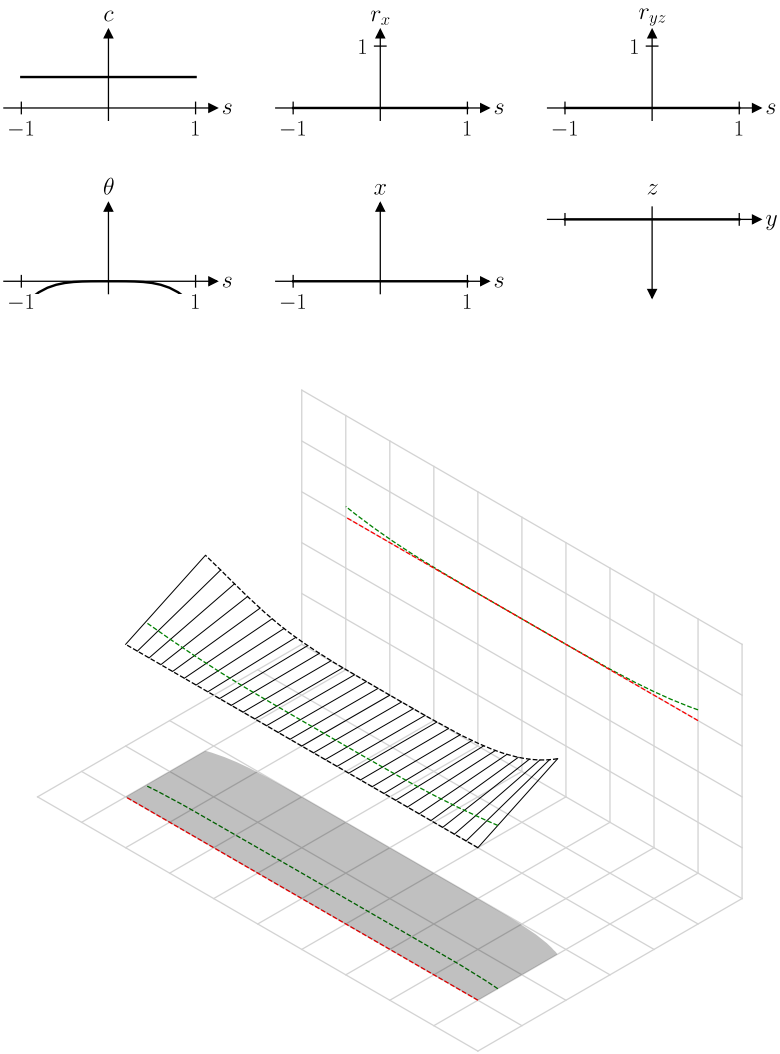

Fig. 3.9: Chord surface of a wing with geometric twist.

3.5.4 Manta ray

The effect of changing the reference positions can be surprising. A great example is a "manta ray" inspired design: each model uses the same piecewise-linear chord distribution and circular $x(s)$, changing only the constant value of r_x. These examples clearly demonstrate the flexibility of the *Simplified model*: four of the six design "curves" are merely constants, and yet they enable significantly nonlinear designs in an intuitive way.

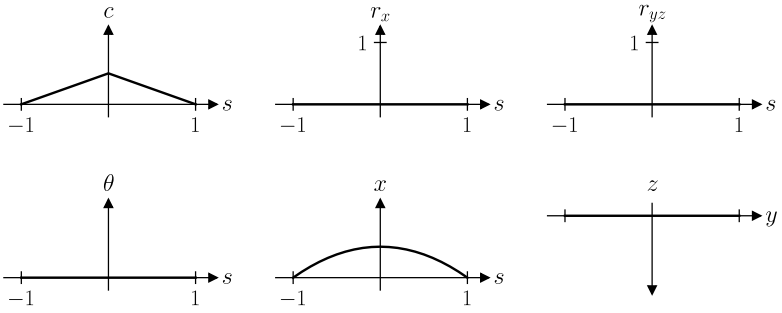

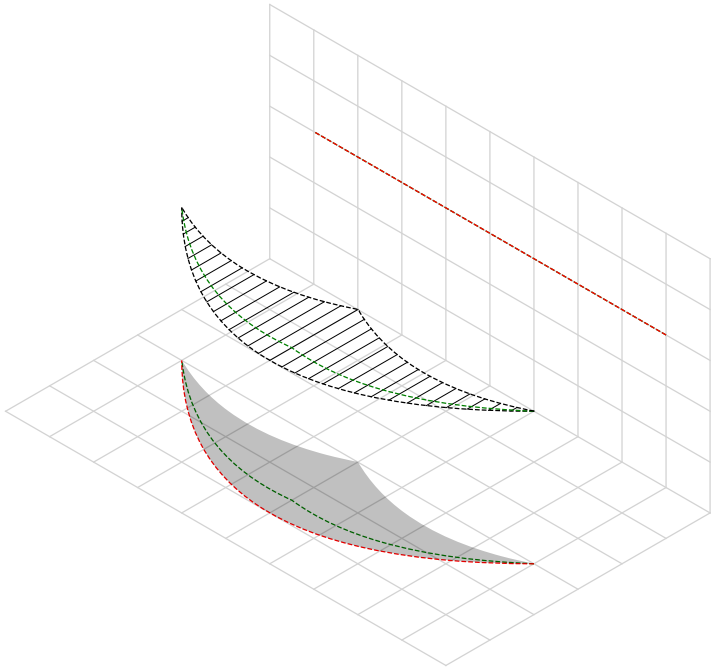

Fig. 3.10: "Manta ray" with $r_x = 0$

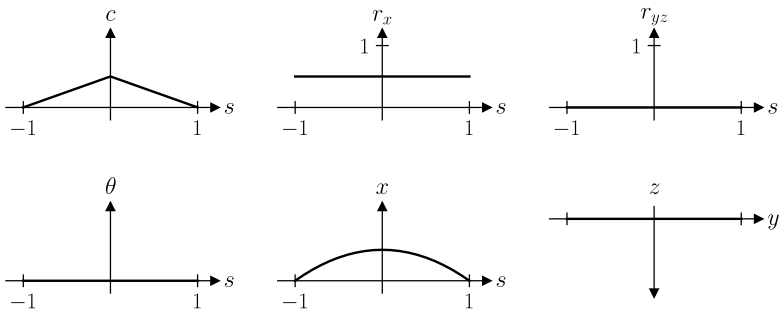

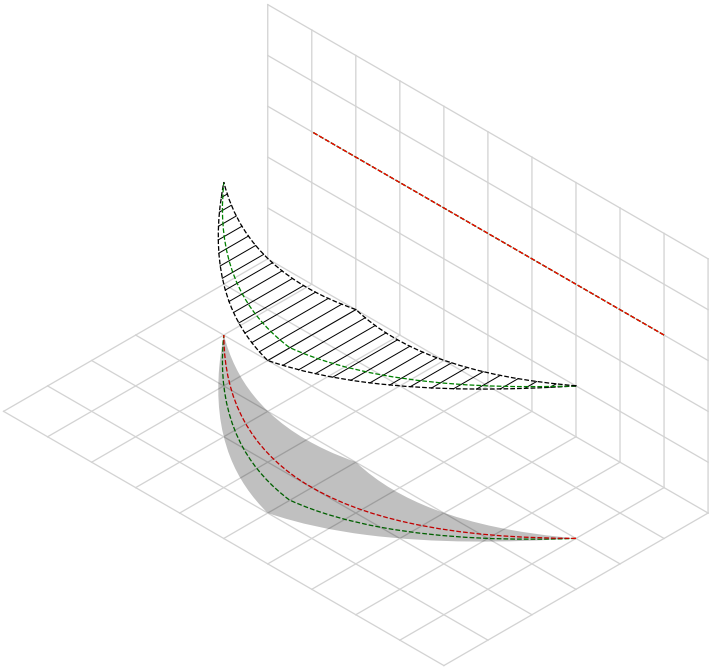

Fig. 3.11: “Manta ray” with $r_x = 0.5$

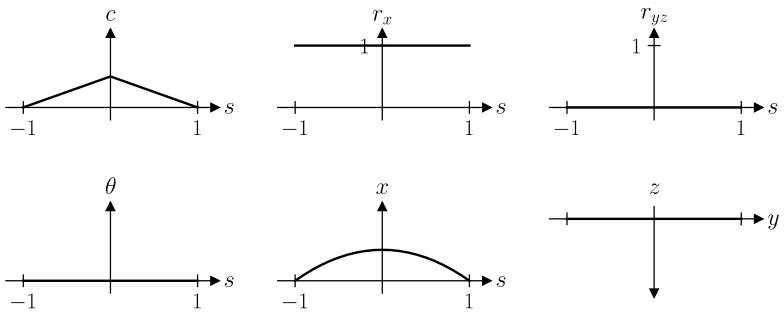

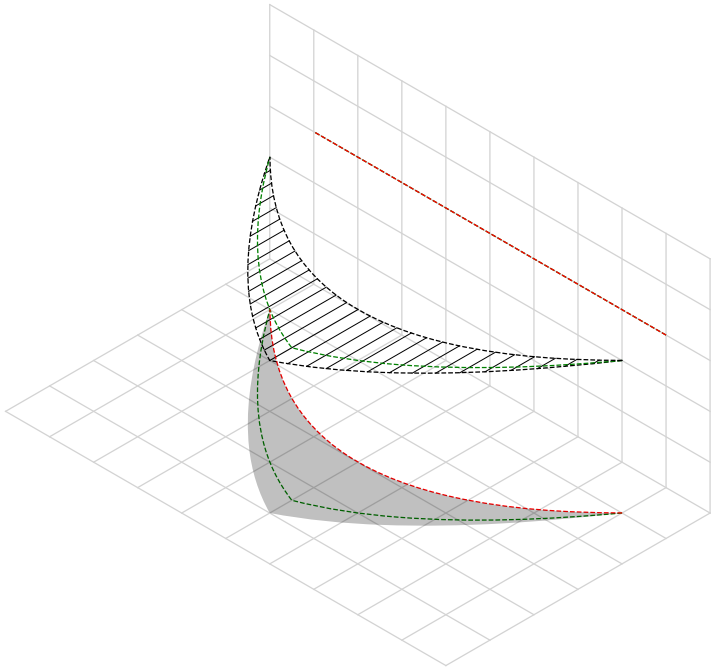

Fig. 3.12: “Manta ray” with $r_x = 1.0$

3.5.5 Parafoil

Lastly, as this project is primarily focused on paragliders, these examples would not be complete without showing how the *Simplified model* allows two simple elliptical functions and $r_x = 0.75$ to easily produce an accurate generalization of a paraglider canopy.

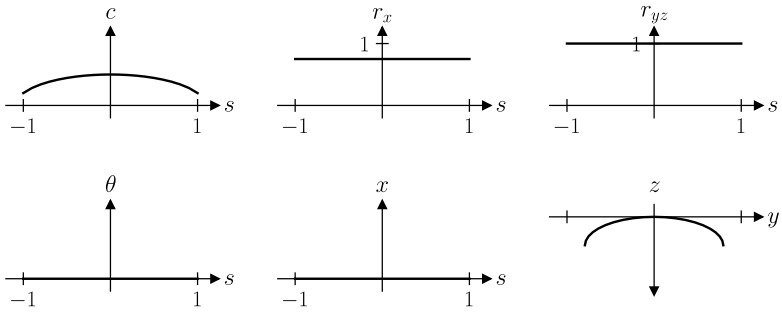

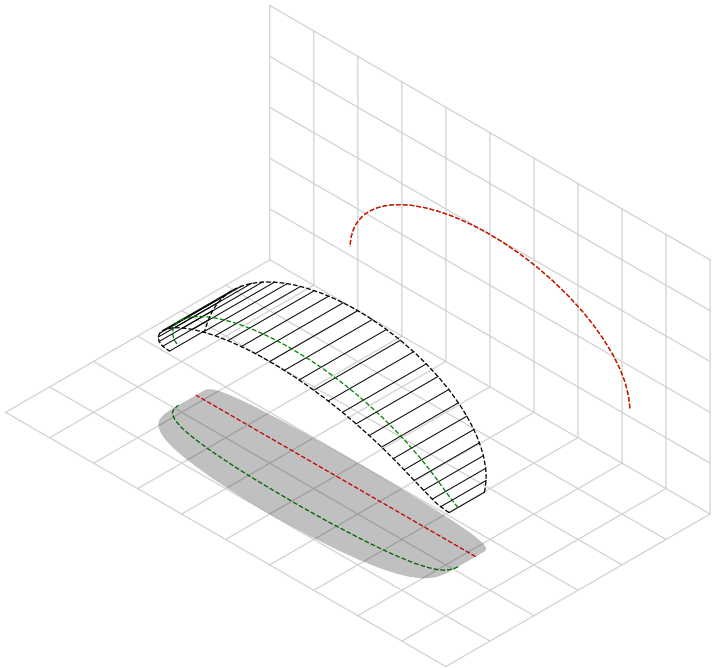

Fig. 3.13: Chord surface of a simple parafoil.

In addition to the surface produced by the section chords, it may be helpful to see the upper and lower profile surfaces produced after assigned every section an airfoil (NACA 23015):

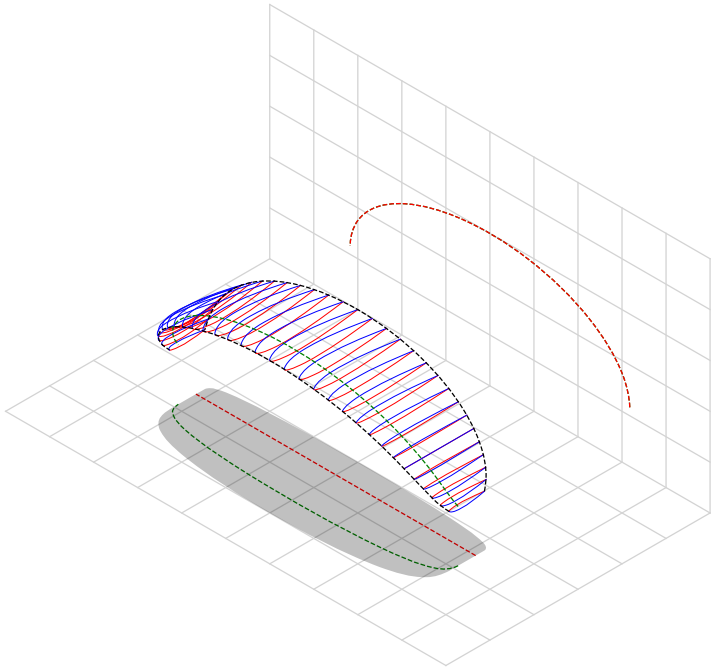

Fig. 3.14: Profile surface of a simple parafoil.

Chapter 4

FOIL AERODYNAMICS

For the purposes of this chapter, an *aerodynamics model* provides the instantaneous forces and moments produced on a foil when it moves relative to air. In a rigorous modeling process the aerodynamic forces and moments would be measured experimentally, either in a wind tunnel or with flight tests, but that rigor is time consuming, expensive, and requires physical possession of the wing. Instead, this paper is concerned with estimating the dynamics of commercial paraglider wings from basic technical specifications, and so it must rely on theoretical methods that predict the flow-field surrounding a foil by combining fundamental equations of fluid behavior with the foil geometry.

This chapter suggests performance criteria for simulating paraglider aerodynamics, and selects a theoretical method capable of simulating those dynamics under the typical flight conditions. It presents a derivation of the method, modifies the method to improve its behavior in the context of flight simulation, and validates the modified method by comparing its predictions against wind tunnel measurements of a representative parafoil model from literature.

4.1 Aerodynamics models

Classical aerodynamics predate the modern computing era, and were forced to prioritize simplifying assumptions that would enable analytical solutions of the governing equations; those assumptions placed heavy restrictions on what geometries could be analyzed and what characteristics of the flow-field must be neglected. These simplifying assumption made the problems tractable in a surprising variety of situations, but — despite their elegance — such analytical solutions are inadequate for analyzing the geometry and flight conditions of a paraglider.

In contrast, modern *computational aerodynamics* [9] solve the equations numerically, relaxing the need for analytical solutions. As a result, modern methods can analyze significantly more complex foil geometries over the entire set of flow-field characteristics. How-

ever, even with modern computers the fluid equations are too difficult to solve in the general case, so simplifying assumptions are still required to produce a tractable system of equations. This modeling process has led to a wide variety aerodynamic models built on different simplifying assumptions regarding the geometry and the characteristics of the flow-field.

4.1.1 Model requirements

The introduction to this paper established a set of *Modeling requirements*, which determine the choice of aerodynamics method. Summarizing those requirements here for convenience, the model must account for the following characteristics:

- Nonlinear geometry
- Viscosity
- Non-uniform wind field (different relative wind angles at different sections)

Where "viscosity" is elaborated as a collection of requirements:

- The model should account for the decreased lift and increased drag due to flow separation across individual wing segments (at least approximately). This requirement is due to paraglider's tendency to fly at relatively high angles of attack, and for individual sections to experience high angles due to the arc anhedral (especially during turns).
- The model must demonstrate graceful accuracy degradation approaching stall (but is not required to model post-stall). The goal is not to simulate with absolute accuracy through stall, but the flight simulator should tolerate brief moments near stall.
- The model should accept empirical corrections to viscous drag to individual wing sections to incorporate experimental wind tunnel results.
- The model should use section-specific Reynolds values (not a wing average) since the sections of a paraglider canopy can vary from 300k to 2M during a turn (thus spanning the transition regime of Reynolds values)

There was also an optional, but desirable, goal that the method should be fast enough for real-time simulations to support rapid iteration during parameter estimation.

4.1.2 Model selection

Despite the wide variety of options for choosing a theoretical aerodynamics model, in practice the *Modeling requirements* makes the selection process rather straightforward. The first requirement — to support nonlinear foil geometries — eliminates the classic LLT. Several authors have developed extensions of the LLT that are able to account for circular arc ([17], [18]), but are unable to model a swept quarter-chord.

The practical answer to nonlinear geometries is to switch to a vortex lattice method or panel method [8], which place the aerodynamic singularities on the nonlinear camber surface, or the profile surface itself, and apply the inviscid flow approximation to reformulate the problem as an instance of Laplace's equation. Unfortunately, the inviscid assumption necessary to produce those solutions violate another of the modeling requirements: the ability to model viscous effects. Although extended models may apply strip theory to incorporate viscous drag coefficients (through lookups based on the estimated section angle of attack or lift coefficient), the inviscid methods fail to provide graceful accuracy degradation near stall. Because the inviscid solutions rely on linear relationships that are assumed to hold indefinitely, they are incapable of capturing the aerodynamic nonlinearities that arise at high angles of attack.

The next level of aerodynamic models are the computational fluid dynamics [9] methods. Instead of limiting the singularities to points on (or inside) the foil, CFD methods simulate the dynamics of the entire volume surrounding the object. In this way they are able to capture the entire array of flow characteristics such as viscosity, turbulence, and compressibility. Unfortunately, CFD methods have the downside of violating another of the modeling requirements: the requirement for speed. The purpose of this project is to enable a user to rapidly iterate the parameters of a model in order to improve the accuracy of a model. Individual CFD simulations at this level are commonly measured in seconds, if not minutes, rendering the fundamentally unsuitable.

Fortunately, there is yet another category, numerical lifting-line methods, which has progressed sufficiently to introduce a method suitable for wings with arbitrary camber, sweep, and dihedral while also supporting (some) viscous effects.

4.2 Phillips' numerical lifting-line

Phillips' numerical lifting-line method (NLLT) [21] is an extension of Prandtl's classic *lifting-line theory* (LLT) to account for the effects of a curved lifting-line.

Unlike the classical LLT, this numerical approach supports the characteristic nonlinear geometry of parafoils by decomposing the foil into discrete wing segments, each with their own scale, position, orientation, and profile. It can also be adapted to non-uniform wind vectors, allowing it to analyze non-uniform, non-longitudinal scenarios involving wind shear and wing rotation.

Unlike pure potential flow solutions, such as traditional vortex lattice and surface panel methods, it is able to approximately account for the effects of viscosity through its use of section coefficients (critical for incorporating viscous drag corrections and approximating flow behavior at high angles of attack).

And unlike full CFD solvers, the implementation is relatively simple, requires minimal manual configuration, and is computationally efficient (a critical point when generating iterated solutions for flight simulation).

4.2.1 Derivation

For the purposes of discussion, the derivation of Phillips' NLLT is briefly repeated here using the notation of this paper. Note that to avoid confusion, this derivation breaks the convention of this paper and instead uses Phillips' convention of a capital $\boldsymbol{V}$ for velocity, and a lowercase $\boldsymbol{v}$ for the induced velocities.

The goal is to establish a system of equations by equating two measures of the aerodynamic force applied to discrete segments of a wing. One uses the 3D vortex lifting law

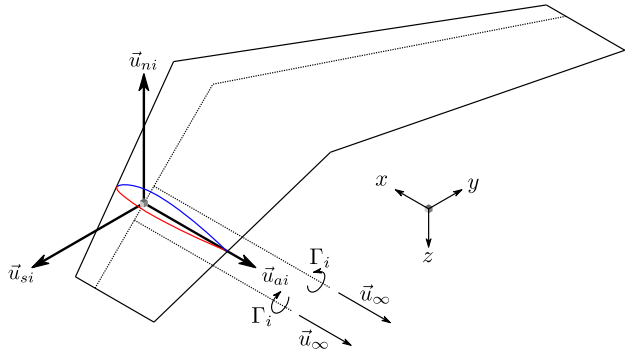

Fig. 4.1: Wing sections for Phillips' method.

(4.1) and the other uses the local section lift coefficients (4.2):

$$\mathbf{d}\boldsymbol{F}_i = \rho\Gamma_i \boldsymbol{V}_i \times \mathrm{d}\boldsymbol{l}_i \tag{4.1}$$

$$\|\mathbf{d}\boldsymbol{F}_i\| = \frac{1}{2}\rho_{\text{air}} \|\boldsymbol{V}_i\|^2 C_{L_i}(\alpha_i, \delta_i) A_i \tag{4.2}$$

The net local velocity $\boldsymbol{V}_i$ at control point i is the sum of the freestream relative wind velocity $\boldsymbol{V}_\infty$ at the control point and the induced velocities from all the other segments:

$$\boldsymbol{V}_i = \boldsymbol{V}_\infty + \sum_{j=1}^{N} \Gamma_j \boldsymbol{v}_{ji} \tag{4.3}$$

where $\boldsymbol{v}_{ji}$ are the velocities induced at control point i by horseshoe vortex j:

$$\boldsymbol{v}_{ji} = \frac{1}{4\pi}\left[\frac{\boldsymbol{u}_\infty \times \boldsymbol{r}_{j_2 i}}{r_{j_2 i}\left(r_{j_2 i} - \boldsymbol{u}_\infty \cdot \boldsymbol{r}_{j_2 i}\right)} + (1-\delta_{ji})\frac{(r_{j_1 i} + r_{j_2 i})(\boldsymbol{r}_{j_1 i} \times \boldsymbol{r}_{j_2 i})}{r_{j_1 i} r_{j_2 i}(r_{j_1 i} r_{j_2 i} + \boldsymbol{r}_{j_1 i} \cdot \boldsymbol{r}_{j_2 i})} - \frac{\boldsymbol{u}_\infty \times \boldsymbol{r}_{j_1 i}}{r_{j_1 i}\left(r_{j_1 i} - \boldsymbol{u}_\infty \cdot \boldsymbol{r}_{j_1 i}\right)}\right] \tag{4.4}$$

and δ_{ji} is the Kronecker delta function:

$$\delta_{ji} \overset{\text{def}}{=} \begin{cases} 1 & i = j \\ 0 & i \neq j \end{cases} \tag{4.5}$$

Solving for the vector of circulation strengths can be approached as a multi-dimensional root-finding problem over f, where f is a vector-valued function of residuals, and the residual for each horseshoe vortex i is the difference between the two measures of section lift, (4.1) and (4.2):

$$f_i(\Gamma_i) = 2\Gamma_i \|\boldsymbol{W}_i\| - \|\boldsymbol{V}_i\|^2 A_i C_{L,i}(\alpha_i, \delta_i) \tag{4.6}$$

where

$$\boldsymbol{W}_i = \boldsymbol{V}_i \times \mathrm{d}\boldsymbol{l}_i \tag{4.7}$$

The set of residuals $f_i(\Gamma_i)$ represent a system of nonlinear equations that can be solved numerically to produce an estimate of the spanwise circulation Γ_i. In order to solve the system, Phillips suggests gradient descent using the system Jacobian $J_{ij} \stackrel{\text{def}}{=} \frac{\partial f_i}{\partial \Gamma_j}$, which expands to:

$$\begin{aligned} J_{ij} = \delta_{ij}\, 2\,\|\boldsymbol{W}_i\| + 2\,\Gamma_i \frac{\boldsymbol{W}_i}{\|\boldsymbol{W}_i\|} \cdot (\boldsymbol{v}_{ji} \times \mathrm{d}\boldsymbol{l}_i) \\ - \|\boldsymbol{V}_i\|^2 A_i \frac{\partial C_{L,i}}{\partial \alpha_i} \frac{V_{a,i}\,(\boldsymbol{v}_{ji} \cdot \boldsymbol{u}_{n,i}) - V_{n,i}\,(\boldsymbol{v}_{ji} \cdot \boldsymbol{u}_{a,i})}{V_{ai}^2 + V_{ni}^2} \\ - 2A_i C_{L,i}(\alpha_i, \delta_i)(\boldsymbol{V}_i \cdot \boldsymbol{v}_{ji}) \end{aligned} \tag{4.8}$$

with the effective wind speed in the normal and chordwise directions

$$\boldsymbol{C}_{f/s_i} = -\begin{bmatrix} | & | & | \\ \boldsymbol{u}_{a,i} & \boldsymbol{u}_{s,i} & \boldsymbol{u}_{n,i} \\ | & | & | \end{bmatrix} \tag{4.9}$$

$$\begin{aligned} V_{a,i} &= \boldsymbol{V}_i \cdot \boldsymbol{u}_{a,i} \\ V_{n,i} &= \boldsymbol{V}_i \cdot \boldsymbol{u}_{n,i} \end{aligned} \tag{4.10}$$

and the *effective local angle of attack* α_i

$$\alpha_i = \arctan\left(\frac{V_{a,i}}{V_{n,i}}\right) \tag{4.11}$$

After solving for the circulation strengths, the 3D vortex lifting law (4.1) is used to compute the inviscid forces at each control point, and the viscous drag and pitching moments are computed as in standard *strip theory* using the effective angle of attack (4.11):

$$\mathrm{d}\boldsymbol{F}_{\text{visc},i} = \frac{1}{2}\rho_{\text{air}}\,\|\boldsymbol{V}_i\|^2\, c_i C_{D,i}\,(\alpha_i, \delta_i)\,\hat{\boldsymbol{V}}_i \tag{4.12}$$

$$\mathrm{d}\boldsymbol{M}_i = -\frac{1}{2}\rho_{\text{air}}\,\|\boldsymbol{V}_i\|^2\, A_i c_i C_{M,i}\,(\alpha_i, \delta_i)\,\boldsymbol{u}_{s,i} \tag{4.13}$$

4.2.2 Modifications

Although the original derivation is suitable for simple, static scenarios, it is inadequate for simulating dynamic conditions that commonly occur during paraglider flights. This section presents a number of modifications to improve the usability, functionality, and numerical stability of the method that greatly extend its applicability.

4.2.2.1 Control point distribution The paper recommends placing the control points using a cosine distribution over the 3D spanwise coordinate y, but that recommendation assumes a predominantly flat wing; cosine spacing generates a poor distribution when the wing tips are nearly vertical, which is common with parafoils. Instead, distributing the control points according to the *section index* s will maintain spacing along the foil's yz-curve regardless of the arc. (Note that although this works well for parafoils, other foil geometries may be better suited to either a different section index, or some nonlinear spacing in s.)

4.2.2.2 Variable Reynolds numbers Lifting-line methods typically assume the section coefficient data is an explicit function of angle of attack α, and possibly some sort of control deflection δ, but assume the coefficients are constant with respect to Reynolds number. For relatively high Reynolds regimes this is reasonable since the airfoil data is essentially constant, but parafoil sections under typical flight conditions experience Reynolds numbers in the range from roughly 150,000 to 3,000,000, spanning the transitional regime where viscous effects can be significant. To verify whether section-local Reynolds numbers have a significant effect on parafoil aerodynamics, the coefficients should be an explicit function of Reynolds number.

4.2.2.3 Non-uniform upstream velocities Phillips' original derivation [21] assumes uniform flow, but [33] relaxes that assumption by replacing the uniform *freestream velocity* V_∞ with the relative *upstream velocity* $V_{rel,i}$ that "may also have contributions from prop-wash or rotations of the lifting surface about the aircraft center of gravity." (Compare Phillips Eq:5 to Hunsaker-Snyder Eq:5.) The result is that (4.3) is replaced with:

$$\boldsymbol{V}_i = \boldsymbol{V}_{rel,i} + \sum_{j=1}^{N} \Gamma_j \boldsymbol{v}_{ji} \tag{4.14}$$

In [33] they are concerned with accounting for propeller wash, but for a parafoil the upstream velocity is simply the local wind velocity at control point i combined with the velocity produced by the control point CP, i rotating about the glider center of mass CM:

$$\boldsymbol{V}_{rel,i} = \boldsymbol{V}_{\infty,i} + \boldsymbol{r}_{CP,i/CM} \times \boldsymbol{\omega}_{b/e} \tag{4.15}$$

This change enables the method to approximately accommodate non-uniform wind conditions, such as from wind shear, turning maneuvers, etc. This flexibility should be used with caution, however; see *Straight-wake assumption* for a discussion.

4.2.2.4 Better solver To solve for the circulation strengths Γ_i, the Phillips paper suggests using *Newtons' method*, which computes the zero of a function via gradient descent. Gradient descent has several practical issues, but the most important problem in this case is that it fails to converge if the gradient goes to zero. For this application, the function under evaluation is the residual error (4.6), and its gradient (4.8) depends on derivatives of the section lift coefficients. When a wing section reaches the angle of attack associated with $C_{L,max}$ the section has stalled, its section lift slope is zero, and gradient descent will fail to converge. Phillips suggests switching to Picard iterations to deal with stalled sections, but it is unclear whether the target function reliably produces fixed points; a simple prototype failed to converge.

An alternative is to use a robust, hybrid root-finding algorithm that uses gradient descent for speed but switches to a line-search method when the gradient goes to zero. The implementation for this project had great success with a modified Powell's method, which

"retains the fast convergence of Newton's method but will also reduce the residual when Newton's method is unreliable" (see the GSL discussion or MINPACK's hybrj documentation for more information). This method not only mitigates the convergence issues near stall, but it is also significantly faster: it does not depend on fixed step sizes (which must be inherently pessimistic to encourage convergence) and is able to use approximate Jacobian updates instead of requiring full Jacobian evaluations at each step.

4.2.2.5 Reference solutions The root-finding algorithm that solves for the circulation strengths requires an initial proposal for the *circulation distribution* $\Gamma(s)$. Poor proposals produce large residual errors that can push Newton iterations into unrecoverable states, so it is preferable to use prior information to predict the true distribution. The original paper suggested solving a linearized version of the equations, but that choice is only suitable for foils with no sweep or dihedral. Another common suggestion from related methods is to assume an elliptical distribution; for most foils, an elliptical circulation distribution is a reasonable guess during straight and steady flight, but it is a poor proposal for scenarios that include non-uniform wind or asymmetric control inputs, such as during flight maneuvers. It is clear that generating suitable proposals for nonlinear geometries under variable flight conditions requires a different approach.

For sequential problems, such as the sequence of states in a flight simulator or the points of a polar curve, an effective solution is to use the solution from the previous iteration as the proposal. Provided the time resolution of the simulation is reasonably small then the state of the aircraft should be similar between each timestep, so the proposal will be very close to the target. An added advantage of using a prior solution is an ability to capture hysteresis effects [38].

4.2.2.6 Clamping section coefficients A major issue with the method is a tendency to produce fictitious "infinite" induced velocities under certain conditions, causing convergence to fail. This tendency increases as the grid resolution is refined, and is most commonly observed at the wing tips, especially during turning maneuvers. The cause is apparent in equation (4.4), where the induced velocities between bound segments increases as the inverse of their separation distance; as the separation distance goes to zero, the induced velocity goes to infinity. In most cases, the induced velocities from the left and right neighbors of a segment mostly cancel, but if the foil has discontinuities (such as at the wingtips, where the outer segment has only an inboard neighbor) then cancellation may be incomplete, leaving a large imbalance. It can also occur due to numerical issues at very fine grid resolutions.

For parafoils the most significant discontinuities are at the wingtips, where the effect of the induced velocity spike is to dramatically overestimate the effective angle of attack. The NLLT relies on accurate section coefficient data, and if that coefficient data is unavailable (such as at high angles of attack) then the numerical routine cannot continue, causing convergence to fail.

Clearly the lack of coefficient data is not a valid reason to abort, since the large induced angle of attack is fictitious. To mitigate the issue when it occurs at the wingtips, assume the true α is less than or equal to the maximum α supported by the coefficient data, and clamp C_L to its value at that maximum α. In the case where the high α is fictitious, the C_L will be incorrect but will at least remain relatively close to the true value, and will allow the simulation to continue. In the case where α is genuinely large, then the unclamped inboard segments will also lack coefficient data and the method will correctly fail.

It is important to note that this is a practical mitigation, not a theoretically-justified solution. The point is not to "fix" the method, the point is to limit the magnitude of the error and allow the simulation to continue with reasonable accuracy. However, despite lacking a theoretical basis, there are several strong justifications:

1. If the outer segment is small, then its contribution to the error is expected to be small. For example, if the outer segment represents the last 5% of the wing span means then

the error from much less than 5% of the total aerodynamic contributions (since the area of that wingtip segment is very small).

2. If the outer segment is small, you wouldn't expect a significant change in alpha from the wingtip to its neighbor, so if the inboard neighbor is in the valid range you can expect that the wingtip alpha is (relatively) close to the valid range.

4.2.3 Limitations

4.2.3.1 Assumes minimal spanwise flow This method argues that the derivation of the 3D vortex lifting law in [39] proves that "the relationship between section lift and section circulation is not affected by flow parallel to the bound vorticity." In other words, it relies on the fact that the 3D vortex lifting law holds even in the presence of spanwise flow. What this does not account for, however, is the effect of spanwise flow on the section coefficients. Wing analysis using section coefficients relies on the assumption that each wing segment acts as a finite segment of an infinite wing, provided the spanwise flow is negligible ([6], p. 356). Although the 3D vortex law holds in the presence of spanwise flow, solving for the circulation strengths using section coefficients does not.

A similar discussion can be found in [38], who apply a similar NLLT to a flat wing with 45° sweep. They acknowledge that although the sweep introduces significant 3D flow-field effects, the method "shows very good agreement" versus experimental measurements. Their success offers some confidence that the effects of spanwise flow may indeed be negligible, but it is unclear whether the effect has more significance once continuous arc anhedral is involved.

4.2.3.2 Straight-wake assumption A common aerodynamic modeling approximation is to assume that vorticity is shed into the wake as a trailing *vortex sheet*; the strength of the shed vorticity varies with the local variation of lift along the span. In a rigorous analysis, the trailing vorticity should follow a curved path ([6], p. 390), but this produces an intractable nonlinear system of equations. Instead, models apply a further simplification known as the

straight-wake assumption: that the trailing *wake vortex sheet* streams straight back from the lifting-line. The straight-wake assumption is an important step in linearizing the system of equations to allow mathematically tractable solutions.

For a discretized method, such as Phillips' or Weissinger's LLT [40], the vortex sheet is lumped into a series of shed vortex filaments whose strength is proportional to the difference in local lift of neighboring segments. Under the straight-wake assumption, the trailing legs of all horseshoe vortices extend from the nodes in straight lines parallel to some *freestream velocity* direction $\boldsymbol{u}_\infty$ (see (4.4)). This is clearly invalid for a rotating wing where a freestream velocity is ambiguous.

Despite this limitation, this project assumes that as long as the rotation rates remain small enough that relative flow angles remain small the method still provides useful approximations. This assumption is made without theoretical justification; instead, this paper relies on the superior aerodynamics knowledge of its sources. First, the use of this method with non-zero rotation is explicitly mentioned in [33]. Also, this assumption is shared with the vortex-lattice model used in AVL [41], although in that method the trailing legs are aligned with the foil x-axis, regardless of freestream flow. In Phillips' method the trailing are aligned to the freestream, which for this work is defined as the local upstream velocity $\boldsymbol{u}_{\infty,0}$ of the central section under the assumption that it minimizes average deviation.

For a related technical discussion that incorporates rotation rates into a vortex lattice method, refer to [8] Sec. 6.5; in particular, Eq. 6.33 for aligning the trailing legs with the x-axis, Eq. 6.37 for accounting by adding it to the flow tangency equations, and Eq. 6.39 for incorporating the rotation rates into the aerodynamic influence coefficients matrix.

4.2.3.3 Reliance on section coefficients A significant limitation of aerodynamic methods based on the theory of *wing sections* their assumption that the section coefficient data is accurate and representative of the flow conditions during a flight. In practice, section coefficient data is notoriously optimistic, relying on idealized geometry, negligible spanwise flow, a uniform flow-field across the segment, steady-state conditions, etc. These assumptions are strong to begin with, and become particularly questionable near stall, especially when

using simulated airfoil data.

Not only do these methods assume the section coefficient data is accurate for each individual section in isolation, they also assume the flow conditions of each section will have a negligible impact on the coefficients of neighboring sections. In reality, development of 3D flow-field conditions such as separation bubbles is significantly impacted by such neighboring sections. Part of the interaction can be captured by the induced velocities, but section coefficients are ultimately incapable of modeling effects such as turbulence, 3D separation bubbles, significant spanwise (or "cross") flow, etc. Such effects seem likely to be even more prominent given the significant arc of a parafoil.

4.2.3.4 No unsteady effects This method produces a steady-state (non-accelerated) solution. It does not include unsteady (time-varying) effects, such as ([8], p. 149):

- Unsteady foil motion
- Unsteady foil deformation
- Spatially-varying or unsteady atmospheric velocity field

Thankfully, the (arguably) most important unsteady effect for the purposes of paraglider simulation under typical flight conditions can be accounted for by the simulator itself; see *Apparent Mass*.

4.2.3.5 Non-unique solutions Gradient descent will find a zero of the residual, but it is not guaranteed to be unique, especially given that the numerical solver relies on tolerances instead of exact solutions. Depending on the initial conditions, the solver may converge to different circulation distributions.

4.2.3.6 Sensitive to initial proposal This method relies on a good proposal (an initial "guess" of the circulation distribution) to encourage convergence while minimizing optimization runtime. The root-finding problem uses the residual error (4.6) which is likely a non-convex function, in which case a global optimization method such as gradient descent is not guaranteed to find the global minimum for a non-convex function, so the solution is sensitive to the starting point (the initial proposal). In practice this issue is not a major problem when the intended use is flight simulation; solutions are generated iteratively, in which case the previous solution is a natural choice for minimizing the initial residual error (see *Reference solutions*). As an added bonus, using the previous solution adds the capability of capturing hysteresis effects [38]; for example, in [42] they discuss a wing that demonstrates hysteresis depending on whether data were generated with increasing versus decreasing alpha. Nevertheless, the fact that the method has a tendency to produce different solutions for different proposals mean the method will exhibit hysteresis effects which may or may not be physically accurate.

4.2.3.7 Unreliable near stall Phillips suggests that this method can be used up to stall "with caution". Closely related to the issues of spanwise flow, the development of stall conditions along a wing has a high likelihood of violating the assumptions used to generate the section coefficients. Worse, the flexible nature of a parafoil will exacerbate the effects of section stall, which cause the profiles to deform and wrinkle even more than normal. Nevertheless, this project attempts to apply the method to "near stall" conditions under the belief that, for the purposes of flight reconstruction, it is preferable to get a low-quality estimate as opposed to no estimate at all. It is vital, however, for the filtering architecture to model the increased uncertainty as sections approach stall conditions.

Chapter 5

COMPONENT MODELS

A paraglider can be modeled as a *system* of three components: a canopy, a harness, and suspension lines that connect the canopy to the harness.

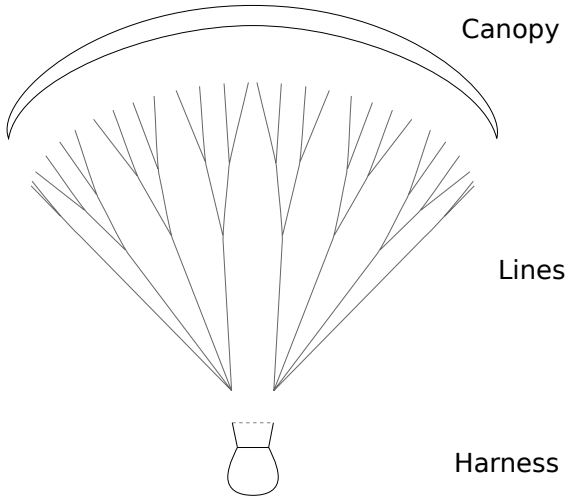

Fig. 5.1: Paraglider component breakdown

Diagram remixed from a Wikipedia contribution by user Mysid.

To compute the dynamics of the composite system, each component model must define three things:

1. Control inputs
2. Inertial properties
3. Resultant force

This chapter develops basic models for each component, favoring simplicity whenever possible. In particular, all models are based on a quasi-rigid body assumption; unlike a true rigid-body model where no component is allowed to move, these models (and their connections) are treated as "instantaneously rigid", where they are allowed specific reconfigurations based on the control inputs (moving the pilot in the harness, or deflecting the trailing edges of the canopy). This may seem like a major oversimplification, but in practice

it works quite well: although nearly every component of a paraglider is made from highly flexible materials, they tend to remain relatively rigid during typical flight conditions.

5.1 Canopy

A paraglider canopy (or *parafoil*) is a kind of ram-air parachute: inflatable lifting surfaces manufactured from nylon sheets with air intakes at the leading edge that pressurize their internal volume. The shape of an inflated parafoil is determined by a combination of surface materials, internal structure, air pressure, and suspension lines. Because the canopy is flexible, pilots can manipulate the suspension lines to change the shape of the canopy, allowing them to control its aerodynamics.

To model a parafoil, it is helpful to think of the canopy as a physical realization of some idealized *foil geometry*. The physical canopy is significantly more complex because it must attempt to create the foil geometry using flexible materials that deform once the canopy is pressurized (as well as meeting requirements such as weight, physical reliability, manufacturability, etc). Modeling the deformations that occur during flight (cell billowing, profile flattening, surface wrinkling, etc) are exceptionally difficult to model without resorting to complete material simulation [36], which is why this project does not consider any deformations other than deflections of the trailing edge due to brake inputs (which are *calculated separately*).

Instead, this model assumes that the foil geometry is an exact representation of the physical canopy, then adds small empirical corrections to account for the most significant error. It models the canopy volume with smooth upper and lower surfaces, whose extents also serve to define the section air intakes. It does not model individual cells, but it does incorporate an estimate of the additional inertia from the internal ribs between each cell. The only deformations included in the model are trailing edge deflections due to pilot control inputs, which are accounted for with precomputed section aerodynamic coefficients; it does not support manipulation via load-bearing lines (used by pilots for maneuvers such as “big ears”, C-riser control, etc) or the *stabilo* lines.

5.1.1 Controls

A paraglider canopy is controlled by changing its shape through manipulation of suspension lines. In theory, any of the suspension lines can be used to alter the positions, orientations, or profiles of its wing sections, but this model only supports trailing edge deflections produced by the lines connected to the left and right brake handles.

When a pilot applies the *brakes*, they generate a continuous deformation along the trailing edge of the canopy. In terms of the individual sections, this results in deformed variants of the undeflected section profiles. Because this canopy model does not perform material simulation, it requires that each variant has been precomputed and assigned a unique *airfoil index* that associates it with a given brake input. The choice of section index has a significant impact on the design of the *suspension line* model, and should be chosen thoughtfully.

A simplistic (but not uncommon) approach is to model the trailing edge deflection as a global rotation about some rotation point, and completely ignore profile deformations. The airfoil index in this case is the deflection angle measured between the deflected and undeflected chords. The rotation point is typically implicit; for example, lifting-line models that assume a fixed quarter-chord are implicitly rotating about the quarter-chord position.

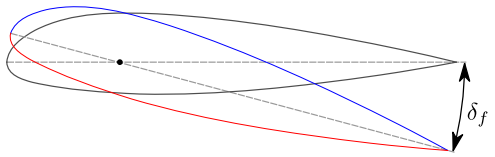

Fig. 5.2: Deflection as a rotation of the entire profile.

By ignoring deformations of the profile geometry this model assumes the shape of the aerodynamic coefficient curves do not change with brake deflections. Instead, the deflection angle δ_f is added directly to the angle of attack, meaning the control input produces a simple translation of the section coefficients. The appeal of this model is the fact that it only requires the section coefficient data from the undeflected profile. Unfortunately, the accuracy of the model degrades rapidly as the deflection angle is increased.

A more accurate model that is extremely common for wings built from rigid materials is to use a discrete *flap* which rotates about a hinge point at some fixed position along the chord:

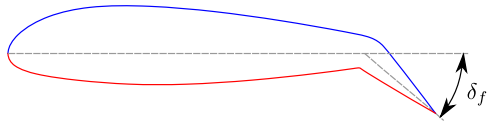

Fig. 5.3: Deflection as a rotation of a rigid flap about a fixed hinge point.

Fixed-hinge flaps are ubiquitous due to their simplicity and acceptable accuracy for rigid wings. Unfortunately, this model is troublesome for flexible wings because there are no fixed hinge points: parafoil edge deflections develop as a variable arc, not a rigid rotation. Also, explicit deflection angles are problematic because parafoil brake inputs cannot control the deflection angles directly; they can only control the downward *deflection distance* δ_d of the trailing edge:

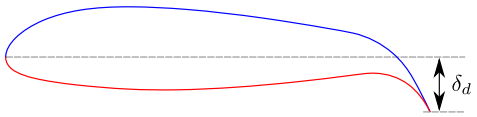

Fig. 5.4: Deflection as a vertical displacement of the trailing edge.

Because airfoils and section coefficients are conventionally normalized to a unit chord, the natural choice of airfoil index for a parafoil is the *normalized deflection distance* $\overline{\delta_d}$, a function of the *deflection distance* δ_d and the *chord length* c:

$$\overline{\delta_d} \stackrel{\text{def}}{=} \frac{\delta_d}{c} \tag{5.1}$$

The normalized deflection distances are unusual in that, although they are control inputs to the canopy aerodynamics model, they are not direct inputs to the system model. Instead, they are computed indirectly using values provided by the *suspension lines* and the *foil geometry* so that the deflection distribution along the span is a function of section index and brake inputs:

$$\overline{\delta_d}\left(s, \delta_{bl}, \delta_{br}\right) = \frac{\delta_d\left(s, \delta_{bl}, \delta_{br}\right)}{c\left(s\right)} \tag{5.2}$$

5.1.2 Inertia

For a parafoil canopy in-flight, the effective inertia is produced by a combination of three different masses: a *solid mass*, from the structural materials, an *air mass*, from the air enclosed in the foil, and an *apparent mass*, from the air surrounding the foil. (Some texts refer to the combination of the solid and enclosed air masses as the *real mass* [24].)

5.1.2.1 Solid mass The *solid mass* is all the surface and structural materials that comprise the canopy. A rigorous model would include the upper and lower surfaces, ribs, half-ribs, v-ribs, horizontal straps, tension rods, tabs (line attachment points), stitching, etc, but for this model the calculation is restricted to the upper and lower surfaces and internal ribs. The internal ribs are assumed to be solid (non-ported), resulting in an overestimate that is somewhat mitigated by the absence of accounting for the other internal structures.

It does, however, account for the extents of the upper and lower surfaces along the section profile. This extent will be used to calculate the inertial properties of the upper and lower surface materials, as well as to calculate empirical viscous correction factors for the section drag coefficients. For this model, the extent of the upper surface and lower surface can be defined using the normalized distance along the section profile, with $-1 \leq r_{\text{lower}} \leq r_{\text{upper}} \leq 1$, with their symmetric spanwise extent controlled by a section index $0 \leq s_{\text{end}} \leq 1$.

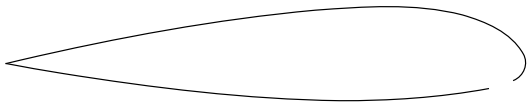

Fig. 5.5: Air intake parameters

Assuming the material densities are uniform, the inertial properties of the materials can be determined by first calculating the total area a and areal inertia matrix $\boldsymbol{J}$ for each surface (using the method in *Area*), then scaling them by the areal densities ρ of each surface. The result is the total masses for the upper surface, lower surface, and internal

ribs:

$$\begin{aligned} m_{\mathrm{u}} &= \rho_{\mathrm{u}} a_{\mathrm{u}} \\ m_{\mathrm{l}} &= \rho_{\mathrm{l}} a_{\mathrm{l}} \\ m_{\mathrm{r}} &= \rho_{\mathrm{r}} a_{\mathrm{r}} \end{aligned} \tag{5.3}$$

And their mass moments of inertia about the canopy origin O:

$$\begin{aligned} \boldsymbol{J}_{\mathrm{u/O}} &= \rho_{\mathrm{u}} \boldsymbol{J}_{a_u/\mathrm{O}} \\ \boldsymbol{J}_{\mathrm{l/O}} &= \rho_{\mathrm{l}} \boldsymbol{J}_{a_l/\mathrm{O}} \\ \boldsymbol{J}_{\mathrm{r/O}} &= \rho_{\mathrm{r}} \boldsymbol{J}_{a_r/\mathrm{O}} \end{aligned} \tag{5.4}$$

In theory the inertial properties are functions of the brake inputs since they alter the distribution of mass, but in practice the effect is negligible. For this project the centroids and moments of inertia for the solid mass are calculated once using the undeflected section profiles.

5.1.2.2 Air mass Although the weight of the air inside the canopy is counteracted by its buoyancy, it still represents significant mass. When the canopy is accelerated the enclosed air is accelerated at the same rate, and must be included in the inertial calculations. (This model neglects surface porosity; although the canopy is porous, and thus constantly receiving an inflow of air through the intakes, in a properly functioning wing the leakage is slow enough that the volume of air can be treated as constant.)

Similar to the surface masses, the internal volume and its unscaled inertia about the canopy origin is easily computed from the *Foil geometry* using the method in *Volume*. Given the internal volume v and the current air density ρ_{air}, the total mass of the enclosed air m_{air} is simply:

$$m_{\mathrm{air}} = \rho_{\mathrm{air}} v \tag{5.5}$$

Similarly, for the inertia matrix of the enclosed air about the canopy origin O:

$$\boldsymbol{J}_{\mathrm{air}/O} = \rho_{\mathrm{air}} \boldsymbol{J}_{\mathrm{v/O}} \tag{5.6}$$

5.1.2.3 Apparent Mass Newton's second law states that the acceleration of an isolated object is proportional to the net force applied to that object:

$$a = \frac{\sum F}{m}$$

This simple rule is sufficient and effective for determining the behavior of isolated objects, but when an object is immersed in a fluid it is longer isolated. When an object moves through a fluid there is an exchange of momentum, and so the momentum of the fluid must be taken into account as well. In fact, it is this exchange of momentum that gives rise to the aerodynamic forces on a wing. The difference is that apparent mass is an unsteady phenomena that is not accounted for by simple aerodynamic models, such as *Phillips' numerical lifting-line*.

In static scenarios, where the vehicle is not changing speed or direction relative to the fluid, this exchange of momentum can be summarized with coefficients that quantify the forces and moments on the wing due to air velocity. But for unsteady flows, where the vehicle is accelerating relative to the fluid, the net force on the vehicle is no longer simply the product of the vehicle's "real" mass and acceleration. Instead, when a net force is applied to an object in a fluid, it will accelerate more slowly than the object would have in isolation, as if the vehicle has increased its mass:

$$a = \frac{\sum F}{m + m_a}$$

This *apparent mass* m_a (or *added mass* [25]) tends to become more significant as the density of the vehicle approaches the density of the fluid. If the density of the vehicle is much greater than the density of the fluid then the effect is often ignored, but for lightweight aircraft the effect can be significant.

Because apparent mass effects are the result of a volume in motion relative to a fluid, its magnitude depends on the volume's shape and the direction of the motion. Unlike the real mass, apparent mass is anisotropic, and the diagonal terms of the apparent inertia matrix are independent. Calculating the apparent mass of an arbitrary geometry is difficult. For a classic discussion of the topic, see [22]. For a more recent discussion of apparent mass in

the context of parafoils, see [23], which used an ellipsoid model to establish a parametric form commonly used in parafoil-payload literature

This paper uses an updated method from [24] which added corrections to the ellipsoid model of [23]. (For a replication of the equations in that method but given in the notation of this paper, see *Apparent mass of a parafoil.*) The method uses several significant simplifying assumptions (the dynamics reference point must lie in the xz-plane, the foil has circular arc, uniform thickness, uniform chord lengths, etc), but the effects of deviations from the method's assumptions are negligible for typical parafoil models.

5.1.3 Resultant force

A method for estimating the canopy aerodynamics was presented *earlier*. An advantage of that method is that it does not assume any particular functional form of the aerodynamic coefficients (linear, polynomial, etc), allowing their definition to use whatever form is convenient. This model uses that flexibility to compose the section coefficients as a two step process:

1. Design a set of airfoils associated with the range of trailing edge deflection, and estimate their aerodynamic coefficients.

2. Apply correction factors to each section to account for physical inaccuracies in the idealized airfoils.

The airfoils are indexed by their normalized deflection distance (5.1), which appears in Phillip' NLLT as the control input δ_i; the indexed airfoils allow the brakes to control the canopy aerodynamics with no modifications to the NLLT. This section index allows each section to provide its own section coefficients, as well as empirical correction factors. One correction factor included in this model, $C_{D,\text{surface}}$, is for "surface roughness" ([43], [30]), and the other, $C_{D,\text{intakes}}$, is for the additional viscous drag due to the air intakes [30]. (See the *demonstration* for an example.) Given the foil geometry and aerodynamic coefficients, the *aerodynamics model* estimates the aerodynamic forces $\boldsymbol{f}_{f,\text{aero},n}$ (4.1) and

moments $\boldsymbol{g}_{f,\text{aero},n}$ (4.13) for the N foil sections.

$$\boldsymbol{f}_{f,\text{weight}} = m_p \boldsymbol{g} \tag{5.7}$$

$$\boldsymbol{f}_{f,\text{aero}} = \sum_{n=1}^{N} \boldsymbol{f}_{f,\text{aero},n} \tag{5.8}$$

$$\boldsymbol{g}_{f/R} = \sum_{n=1}^{N} \left(\boldsymbol{r}_{CP_n/R} \times \boldsymbol{f}_{f,\text{aero},n} \right) + \sum_{n=1}^{N} \boldsymbol{g}_{f,\text{aero},n} + \boldsymbol{r}_{S/R} \times \boldsymbol{f}_{f,\text{weight}} \tag{5.9}$$

5.1.4 Parameter summary

In addition to the design curves that define the *Foil geometry*, the physical canopy model requires additional information about physical details associated with that geometry:

r_{upper}	Profile extent of the upper surface
r_{lower}	Profile extent of the lower surface
s_{end}	Section index where air intakes end
ρ_u	Areal density of the upper surface material
ρ_r	Areal density of the internal rib material
ρ_l	Areal density of the lower surface material
N_{cells}	Number of internal cells
$C_{D,\text{intakes}}$	Drag coefficient due to air intakes
$C_{D,\text{surface}}$	Drag coefficient due to surface characterstics

(5.10)

5.2 Suspension lines

The suspension lines connect the canopy to the harness and pilot. The lines are conventionally grouped into load-bearing sets (labeled A/B/C/D, depending on their relative positions on the section chords), brake lines (that produce the trailing edge deflections), and *stabilo* lines (that assist in preventing the wing tips from curling into a dangerous *cravat*). Starting from the canopy, the lines progressively attach together in a *cascade* that terminates at two

risers which connect the lines to the harness. The lines are responsible for producing the arc of the canopy, suspending the harness at some position relative to the canopy, and allowing the pilot to manipulate the shape of the canopy.

For rigorous models the line geometry is a major factor in wing performance, but for this project a fully-specified suspension line model would be both tedious and redundant. It would be tedious because it would require the lengths of every segment of every line, and it would be (mostly) redundant because the *canopy model* is a quasi-rigid body whose *arc* is already defined by the yz-curve of the idealized foil geometry. As a result, the suspension lines can only affect the riser position and trailing edge deflections, so this model can reasonably use simple approximations that do not depend on an explicit line geometry.

5.2.1 Controls

The suspension lines provide two primary methods of controlling the paraglider system: through brakes, which change the canopy aerodynamics, and the accelerator, which repositions the payload underneath the canopy.

5.2.1.1 Brakes A parafoil canopy can be manipulated by pulling on any of its many suspension lines, but two of the lines in particular are dedicated to slowing the wing or controlling its turning motion. Known as the *brakes* or *toggles*, these controls induce downward trailing edge deflections (see Fig. 5.4) along each half of the canopy, increasing drag on that side of the wing. Symmetric deflections slow the wing down, and asymmetric deflections cause the wing to turn.

A physically accurate model of the deflection distribution would need to model the length and angle of every line and how the angles deform during braking maneuvers. Because the line geometry was not a focus for this project, an approximation is used instead.

First, observe that as brakes are progressively applied the deflections will typically start near the middle and radiate towards the wing root and tip as the brake magnitude is increased. For small brake inputs the deflections are zero near the wing root and tip, but for large brake inputs even those sections experience deflections.

Fig. 5.6: Asymmetric brake deflection.

Photograph by Frédéric Bonifas, distributed under a CC-BY-SA 3.0 license.

Fig. 5.7: Symmetric brake deflection.

Photograph by Wikimedia contributor "PiRK" under a CC-BY-SA 3.0 license.

To approximate this behavior, start by assuming the deflection distances from each individual brake input are symmetric around some peak near the middle of each semispan and vary as a quartic function $q(p)$. Define the polynomial coefficients such that the function value and slope are zero at $p = 0$ and $p = 1$ and a peak at $p = 0.5$. The result is a quartic that is symmetric about $p = 0.5$ with a peak magnitude of 1.

$$q(p) = \begin{cases} 16p^4 - 32p^3 + 16p^2 & 0 \leq p \leq 1 \\ 0 & \text{else} \end{cases} \tag{5.11}$$

Next define two variables for the section indices near the canopy root and tip that control the start and stop points of the deflection. Representing the start and stop positions as variables allows modeling how the deflection distribution changes with the brake inputs.

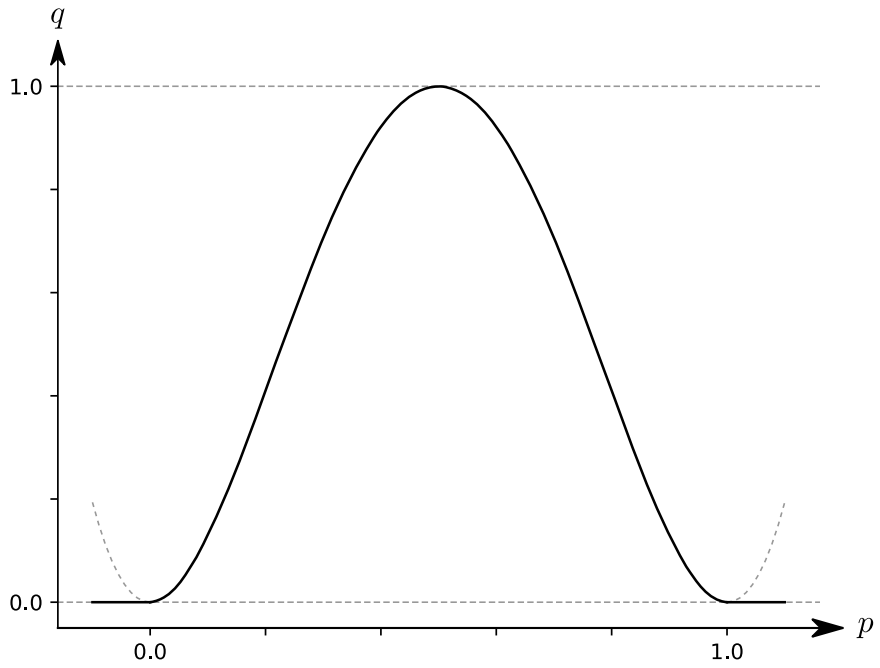

Fig. 5.8: Truncated quartic distribution

For both s_{start} and s_{stop}, define their values when $\delta_{br} = 0$ and $\delta_{br} = 1$. Then, using linear interpolation as a function of brake input:

$$\begin{aligned} s_{\text{start}} &= s_{\text{start},0} + (s_{\text{start},1} - s_{\text{start},0})\,\delta_b \\ s_{\text{stop}} &= s_{\text{stop},0} + (s_{\text{stop},1} - s_{\text{stop},0})\,\delta_b \end{aligned} \tag{5.12}$$

The start and stop points can be used to map the section indices s into the domain of the quartic p, such that $s = s_{\text{start}} \rightarrow p = 0$ and $s = s_{\text{stop}} \rightarrow p = 1$:

$$p(s) = \frac{s - s_{\text{start}}}{s_{\text{stop}} - s_{\text{start}}} \tag{5.13}$$

The quartic output for each brake is unit magnitude, which should be scaled by the brake input. Summing the two scaled outputs represent the fraction of maximum brake deflection distance over the entire span. The maximum brake deflection distance is a constraint set by the suspension line model parameter κ_b, the maximum length that the model will allow the pilot to pull the brake line (although on a physical wing there isn't a clear limit to how far the brakes can be pulled).

Finally, the total brake deflection distance is the sum of contributions from left and right brake:

$$\delta_d(s, \delta_{bl}, \delta_{br}) = (\delta_{bl} \cdot q(p(-s)) + \delta_{br} \cdot q(p(s))) \cdot \kappa_b \tag{5.14}$$

A feature of this design is that setting $s_{\text{start},1} < 0$ allows deep brake inputs to deflect the opposing semispan, and $s_{\text{stop},1} > 1$ allows deflections at the wing tips, as shown in Fig. 5.9.

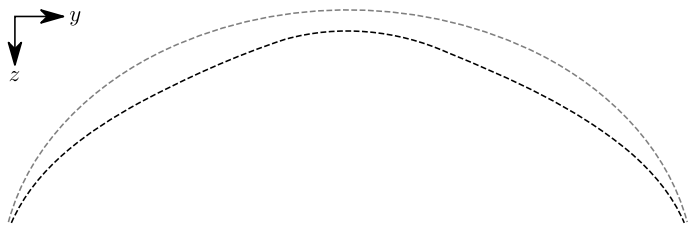

Fig. 5.9: Quartic brake deflections, $\delta_{bl} = 1.00$ and $\delta_{br} = 1.0$

Together with the *Foil geometry*, the absolute brake deflection distances can be used to compute each section's *airfoil index* (5.1).

5.2.1.2 Accelerator Paragliders are not powered aircraft, but pilots can increase their airspeed by adjusting how the payload is positioned relative to the canopy. The *accelerator* or *speed bar* is positioned under the pilot's feet, and by pushing out they can shift the riser position RM forward and up. The canopy pitching angle, angle of attack, and airspeed must adjust to the new equilibrium, changing both the airspeed and the glide ratio.

The goal is to model how the riser position changes as a function of the accelerator control input $0 \leq \delta_a \leq 1$.

For notational simplicity, define $\overline{A}$ and $\overline{C}$ as the lengths of the lines connecting them to the riser midpoint RM:

$$\overline{A} \stackrel{\text{def}}{=} \left\| \boldsymbol{r}_{A/RM} \right\|$$
$$\overline{C} \stackrel{\text{def}}{=} \left\| \boldsymbol{r}_{C/RM} \right\|$$

The default lengths of the lines are defined by two pairs of design parameters. First, the default position of the riser midpoint RM is defined with κ_x and κ_z; this is the position of

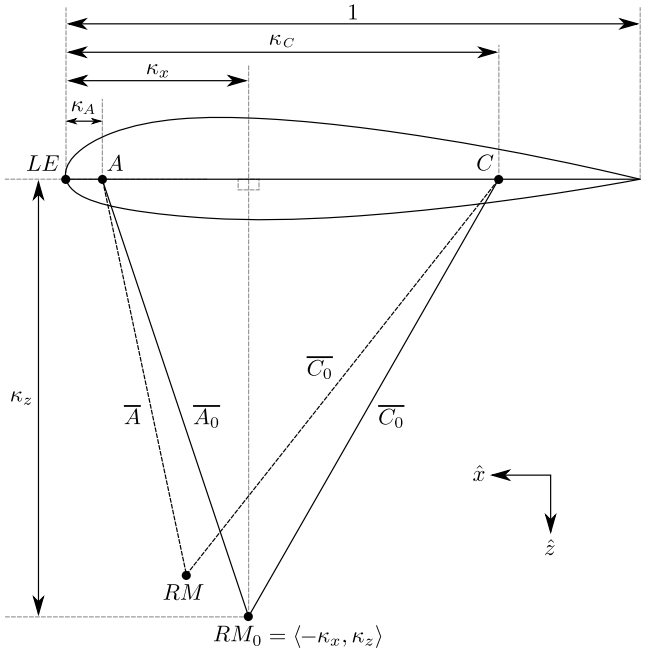

Fig. 5.10: Paraglider wing accelerator geometry.

RM when $\delta_a = 0$. Second, two connection points along the canopy root chord are defined with κ_A and κ_C; connecting lines from these points are the physical means by which RM is positioned underneath the canopy. The A lines connect near the front of the wing, and are variable length; the pilot can use the *accelerator* to shorten the lengths of these lines. The C lines connect towards the rear of the canopy, and are fixed length.

Geometrically, shortening $\overline{A}$ will move RM forward while rotating the C lines. Aerodynamically, shortening $\overline{A}$ effectively rotates the canopy pitch down about the point C, decreasing the global angle of incidence of the canopy; decreasing the angle of incidence decreases lift, and the wing must accelerate to reestablish equilibrium.

A fifth design parameter, the *accelerator length* κ_a, is required to define the maximum length change produced by the accelerator; this is the maximum length that $\overline{A}$ can be decreased. This value is limited by the physical geometry of the pulleys that give the pilot the leverage to pull the canopy into its new position. The pilot uses the *accelerator control input* δ_a, a value between 0 and 1, to specify the total decrease in $\overline{A}$:

$$\overline{A}(\delta_a) = \overline{A_0} - \delta_a \kappa_a \tag{5.15}$$

For deriving the basic geometric relations, it is convenient to normalize all the design parameters by the central chord. This avoids the extra terms in the derivation and allows a wing design to scale naturally with the canopy.

The goal is to use the physical geometry, where the risers position is determined by $\overline{A}$ and $\overline{C}$, to define the position of RM a function of δ_a. The first step is to determine the default line lengths by setting $\delta_a = 0$ and applying the Pythagorean theorem:

$$\begin{aligned} \overline{A_0} &= \sqrt{\kappa_z^2 + (\kappa_x - \kappa_A)^2} \\ \overline{C_0} &= \sqrt{\kappa_z^2 + (\kappa_C - \kappa_x)^2} \end{aligned} \tag{5.16}$$

In the general case, the line lengths are functions of δ_a:

$$\begin{aligned} \overline{A}(\delta_a)^2 &= RM_z^2 + (RM_x - \kappa_A)^2 \\ \overline{C}(\delta_a)^2 &= RM_z^2 + (\kappa_C - RM_x)^2 = \overline{C_0}^2 \end{aligned} \tag{5.17}$$

Where $\overline{C} \equiv \overline{C_0}$ due to the physical constraint that the length of the C lines are constant.

Subtract the two equations in (5.17):

$$\overline{A}(\delta_a)^2 - \overline{C_0}^2 = (RM_x - \kappa_A)^2 - (\kappa_C - RM_x)^2$$

Finally, substitute (5.15) and solve for RM_x and RM_z as functions of δ_a:

$$\begin{aligned} RM_x(\delta_a) &= \frac{\left(\overline{A_0} - \delta_a \kappa_a\right)^2 - \overline{C_0}^2 - \kappa_A^2 + \kappa_C^2}{2\left(\kappa_C - \kappa_A\right)} \\ RM_z(\delta_a) &= \sqrt{\overline{C_0}^2 - \left(\kappa_C - RM_x(\delta_a)\right)^2} \end{aligned} \tag{5.18}$$

The final position of RM with respect to the leading edge (which is also the origin of the canopy coordinate system), scaled by the length of the central chord c_0 of the wing, is then:

$$\boldsymbol{r}^b_{RM/LE}(\delta_a) = c_0 \cdot \langle -RM_x(\delta_a), 0, RM_z(\delta_a) \rangle \tag{5.19}$$

Where RM_x was negated since the wing x-axis is positive forward.

5.2.2 Inertia

This simplistic model assumes the inertia of the lines is negligible compared to that of the canopy; in particular, inaccuracies in the simplified canopy inertia are more significant than the line inertia, so this model simply defines the translational and rotation inertia as zero.

5.2.3 Resultant force

Although the lines are nearly invisible compared to the rest of the wing, they contribute a significant amount of aerodynamic drag. Because the total system drag of a paraglider is relatively small, even a small increase can have a large impact on sensitive characteristics such as glide ratio; in fact, paraglider suspension lines contribute upwards of 20% of the total paraglider system drag ([30], [20]), and should not be neglected.

This model does not provide an explicit line geometry, so it can't compute the true line area distribution. Instead, it lumps the entire length of the lines into configurable control points; for example, given the total line length and average line diameter, the line area can be lumped into singularities such as the centroid of line area for each semispan. As with other similar designs [20], this model treats the drag as isotropic (because the operating ranges of alpha and beta are so small the line drag is effectively constant, and what little force exists along the z-axis is negligible compared to the lift of the canopy). Given the total area S_{lines} represented by each singularity the total aerodynamic drag at some control point L can be calculated as in [20] or [30]:

$$S_l = \kappa_L \kappa_d \tag{5.20}$$

$$\boldsymbol{f}_{l,\text{aero},n} = \frac{1}{2}\rho_{\text{air}} \left\|\boldsymbol{v}_{W/L_n}\right\|^2 S_l C_{d,l,n} \hat{\boldsymbol{v}}_{W/L_n} \tag{5.21}$$

$$\boldsymbol{f}_{l,\text{aero}} = \frac{1}{N}\sum_{n=1}^{N} \boldsymbol{f}_{l,\text{aero},n} \tag{5.22}$$

$$\boldsymbol{g}_{l/R} = \frac{1}{N}\sum_{n=1}^{N} \boldsymbol{r}_{CP_n/R} \times \boldsymbol{f}_{l,\text{aero},n} \tag{5.23}$$

5.2.4 Parameter summary

For the harness position:

κ_A	Chord ratio to the A lines
κ_C	Chord ratio to the C lines
κ_x	Chord ratio to the x-coordinate of the riser midpoint
κ_z	Chord ratio to the z-coordinate of the riser midpoint
κ_a	Accelerator line length

(5.24)

For the brakes:

$s_{\text{start},0}, s_{\text{start},1}$	Section indices where deflections begin for $\delta_b \in \{0, 1\}$
$s_{\text{stop},0}, s_{\text{stop},1}$	Section indices where deflections end for $\delta_b \in \{0, 1\}$
κ_b	Maximum trailing edge deflection distance

(5.25)

For the aerodynamics:

κ_L	Total line length
κ_d	Average line diameter
$\boldsymbol{r}_{CP_n/R}$	Position of lumped control point n
$C_{d,l,n}$	Line drag coefficient for control point n

(5.26)

5.3 Harness

A paraglider harness is the seat for the pilot, which is suspended from the risers. Safety straps over the legs and chest ensure the pilot cannot fall from the harness in turbulent conditions or during unsteady maneuvers. A tensioning strap in front of the pilot's chest controls the horizontal riser separation distance, which allows the pilot to adjust the balance between stability (sensitivity to turbulence) and wing responsiveness to weight shift control. In addition to giving the pilot a safe place to sit, the harness also provides places to store the pilot's gear, a pouch to contain the emergency reserve parachute, and optional padding to protect the pilot in the event of a crash.

Instead of attempting to capture all the geometric irregularities of paraglider harnesses, this model calls upon a time-honored solution from physics: it considers the harness as a sphere. Moreover, the pilot, gear, and reserve parachute are accounted for by simply adding their masses to the mass of the harness. The harness, pilot, and gear are collectively referred to as the *payload*.

5.3.1 Controls

Paraglider harnesses allow pilots to shift their weight left and right, causing an imbalanced load on each semispan. (For a real wing this maneuver also causes a vertical shearing stress along the center of the foil, but due to the rigid body assumption of the canopy model this deformation will be neglected.) The weight imbalance causes the canopy to roll towards the shifted mass, resulting in a gentle turn in the desired direction. Although the turn rate is less than can be produced by the brakes, this maneuver causes less drag and is preferred (when suitable) for its aerodynamic efficiency.

The movement of the pilot can be arguably described as occurring inside the volume of the harness, so *weight shift* control can be modeled as a displacement of the payload center of mass P. Given that the pilot can only shift a limited distance κ_w in either direction, a natural choice of control input is $-1 \leq \delta_w \leq 1$. With the harness initially centered in the canopy xz-plane, the displacement due to weight shift control is $\Delta y = \delta_w \kappa_w$. The displacement of the payload center of mass produces a moment on the risers that rolls the wing and induces the turn.

Defining the riser midpoint RM as the origin the harness-local coordinate system, the position of the displaced center of mass is then:

$$\boldsymbol{r}_{P/RM} = \bar{\boldsymbol{r}}_{P/RM} + \langle 0, \delta_w \kappa_w, 0 \rangle \tag{5.27}$$

5.3.2 Inertia

As in [44] (and similarly in [20]), the payload is modeled as a solid sphere of uniform density. With a total mass m_p, center of mass P, and projected surface area S_p, the moment of inertia about the payload center of mass is simply:

$$\boldsymbol{J}_{p/P} = \begin{bmatrix} J_{xx} & 0 & 0 \\ 0 & J_{yy} & 0 \\ 0 & 0 & J_{zz} \end{bmatrix}$$

where

$$J_{xx} = J_{yy} = J_{zz} = \frac{2}{5} m_p r_p^2 = \frac{2}{5} \frac{m_p S_p}{\pi}$$

5.3.3 Resultant force

Harness drag coefficients were studied experimentally in [44]. The author measured several harness models in a wind tunnel and converted the results into aerodynamic coefficients normalized by the cross-sectional area of the sphere. For a more sophisticated approach the coefficient can be adjusted to account (approximately) for angle of attack and Reynolds number [20], but this model simply treats the drag coefficient as a constant.

$$\boldsymbol{f}_{p,\text{weight}} = m_p \boldsymbol{g} \tag{5.28}$$

$$\boldsymbol{f}_{p,\text{aero}} = \frac{1}{2} \rho_{\text{air}} \left\| \boldsymbol{v}_{W/P} \right\|^2 S_p C_{D,p} \hat{\boldsymbol{v}}_{W/P} \tag{5.29}$$

$$\boldsymbol{g}_{p/R} = \boldsymbol{r}_{CP/R} \times \boldsymbol{f}_{p,\text{aero}} + \boldsymbol{r}_{P/R} \times \boldsymbol{f}_{p,\text{weight}} \tag{5.30}$$

Note that the spherical nature of the model implies isotropic drag. Although this is clearly a poor assumption for such a significantly non-spherical object, the fact that the wind is rarely more than 15 degrees off the x-axis means the such a "naive" drag coefficient will remain fairly accurate over the typical range of operation (regardless of the poor geometric accuracy). This assumption also has the downside that it will never produce an aerodynamic moment about the payload center of mass, but in the absence of experimental data on the magnitude of the missing moment, this model continues to ignore it.

5.3.4 Parameter summary

m_p	Total payload mass
$\bar{\boldsymbol{r}}_{P/RM}$	Payload center of mass default position
κ_w	Maximum weight shift distance
S_p	Projected payload area
$C_{d,p}$	Payload drag coefficient

Chapter 6

SYSTEM DYNAMICS

This chapter combines the individual *component models* into composite *system dynamics* models. In this paper, a *system dynamics* model is a set of derivatives that define the translational and angular acceleration of groups of components that represent an aircraft, specified using a coordinate system attached to the aircraft. Developing a system model can be roughly described as a sequence of steps:

1. Choose a set of components to represent the aircraft
2. Characterize their connections
3. Choose a dynamics reference point for the composite system
4. Develop the system of equations for the accelerations

6.1 Components

The previous chapter defined component models for the canopy, suspension lines, and harness; in the system models, these are lumped into two quasi-rigid-body groups called the *body* and the *payload*. The *body* of the glider is the combination of canopy and suspension lines. The *payload* includes the harness, pilot, and their gear (in this simplified model, the pilot and their gear are treated as additional masses that are added to the mass of the harness).

These models are quasi-rigid because the dynamics equations will only consider their instantaneous configurations when calculating their accelerations; conservation of momentum requires accounting for redistributions of mass, but doing so would require inertia derivatives as functions of time derivatives of the control input (such as weight shift, accelerator, etc), which would significantly complicate the model. Because the redistributions of mass are relatively small for typical scenarios, these models assume the affect of violating conservation of momentum is negligible.

It is important to note that the unfortunately ambiguous terminology of *body* is deliberate. The paraglider community typically refers to the combination of canopy and lines as a *paraglider wing*, but the "body" convention improves consistency with existing *parafoil-payload* literature (which in turn inherited the term from conventional aeronautics literature). Some texts prefer the term *parafoil*, but having the same prefix p for both *parafoil* and *payload* makes subscripting the variables unnecessarily difficult. Similarly, using "wing" would be preferred in this context, but subscripting with w causes confusion when discussing wind vectors. Referring to whatever group of components include the canopy as the *body* was a compromise chosen for consistency with existing literature.

6.2 Connections

Next, the system model must characterize the connection between the body and payload. In literature, parafoil-payload models are commonly categorized by their *degrees-of-freedom* (DoF): the total number of dimensions in which the components of the system are free to move. The body has 3-DoF for translational motion and another 3-DoF for rotational motion, and if the payload is allowed to translate or rotate relative to the body, those additional DoF are added to the total DoF of the system model. For example, in a 6-DoF model, the body and payload are connected as a single rigid body, with no relative motion between them.

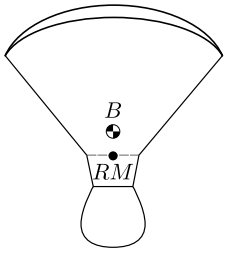

Fig. 6.1: Diagram for a 6-DoF model.

For typical paragliding flight maneuvers, assuming a fixed payload orientation is reasonably accurate, but with one significant failing: although the relative roll and twist are typically negligible, relative pitch about the riser connections is very common, even during

static glides. Friction at the riser carabiners (and aerodynamic drag, to a lesser extent) dampen pitching oscillations, but the payload is otherwise free to pitch as necessary to maintain equilibrium. Assuming a fixed relative pitch angle introduces a fictitious pitching moment that disturbs the equilibrium conditions of the wing and artificially dampens the pitching dynamics during maneuvers. To mitigate that issue, the obvious solution is to add an additional DoF, but for demonstration purposes it is simpler to define a full 9-DoF model, where the body and payload are connected at the *riser midpoint RM*. The connection is modeled as a spring-damper system, which produces an internal force $\boldsymbol{F}_R$ and moment $\boldsymbol{M}_R$:

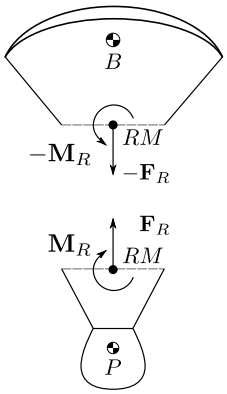

Fig. 6.2: Diagram for a 9-DoF model with internal forces.

6.3 Reference point

Each dynamics model must choose a reference point about which the moments and angular inertia are calculated. A common choice for conventional aircraft is the center of real mass because it decouples the translational and angular dynamics of isolated objects. For a paraglider, however, this is not possible: paragliders are sensitive to apparent mass, which depends on the direction of motion, so there is no "center" that decouples the translational and rotational terms of the apparent inertia matrix [24]. Because the system matrix cannot be diagonalized there is no advantage in choosing the center of real mass. Instead, the reference point can be chosen such that it simplifies other calculations.

In particular, the *method* to estimate the apparent inertia matrix requires that the reference point lies in the xz-plane of the canopy. Two natural choices in that plane are the leading edge of the central section, or the midpoint between the two risers. The *riser midpoint RM* has the advantage that is a fixed point in both the body and payload coordinate systems, which means it does not depend on the relative position or orientation of the payload with respect to the body. (This choice simplifies the equations for the *9-DoF model* while maintaining consistency with the *6-DoF model*.)

6.4 System inputs

The inputs $\boldsymbol{u}$ to the system model the control inputs for each component (with the exception of the trailing edge *deflection distances* $\delta_d(s)$ which are computed internally using the suspension lines and foil geometry models), the wind velocity $\boldsymbol{v}_{W/e}$, air density ρ_{air}, and the gravity vector $\boldsymbol{g}$.

$$\boldsymbol{u} = \left\{ \delta_a, \delta_{bl}, \delta_{br}, \delta_w, \boldsymbol{v}^b_{W/e}, \rho_{\text{air}}, \boldsymbol{g}^b, \right\} \tag{6.1}$$

Here the wind field is assumed to be uniform so the wind velocity at every control point is defined by a single, constant vector, but for non-uniform wind fields there will be a unique wind vector for each aerodynamic control point.

6.5 Equations of motion

The equations of motion are developed by solving for the derivatives of translational momentum ${}^e\dot{\boldsymbol{p}} = \sum \boldsymbol{F} = m\dot{\boldsymbol{v}}$ and angular momentum ${}^e\dot{\boldsymbol{h}} = \sum \boldsymbol{M} = \boldsymbol{J}\dot{\boldsymbol{\omega}}$ for each group of components [11]. In addition to requiring the forces, moments, and inertia matrices for each component, each system model must choose a dynamics reference point and whether to account for the affects of *apparent mass*. The *appendix* includes derivations demonstrating different choices for several each model.

For the 6-DoF model, the most complete is *Model 6a* which accounts for the effects of apparent mass, while *Model 6b* and *Model 6c* have the advantage of simplicity (making them

easier to implement and useful for validating implementations of more complex models). The derivation produces a system of equations (2.13) that can be solved for the two vector derivatives that describe the accelerations of the body relative to the earth frame $\mathcal{F}_e$ taken with respect to the body frame $\mathcal{F}_b$:

$$\begin{aligned} {}^{b}\dot{\boldsymbol{v}}_{RM/e} &\quad \text{translational acceleration of the riser midpoint } RM \\ {}^{b}\dot{\boldsymbol{\omega}}_{b/e} &\quad \text{angular acceleration of the body} \end{aligned} \tag{6.2}$$

Similarly, for the 9-DoF model, *Model 9a* also develops a complete system of equations (2.32) that account for apparent mass of the canopy, but with the addition of a separate angular acceleration for the payload with respect to the payload frame $\mathcal{F}_p$:

$$\begin{aligned} {}^{b}\dot{\boldsymbol{v}}_{RM/e} &\quad \text{translational acceleration of the riser midpoint } RM \\ {}^{b}\dot{\boldsymbol{\omega}}_{b/e} &\quad \text{angular acceleration of the body} \\ {}^{p}\dot{\boldsymbol{\omega}}_{p/e} &\quad \text{angular acceleration of the payload} \end{aligned} \tag{6.3}$$

Chapter 7

STATE DYNAMICS

The *System dynamics* defined the instantaneous accelerations of the aircraft in terms of local reference frames traveling with the aircraft. To record the behavior of an aircraft over time, a set of variables must be chosen to encode the *state* of the system relative to some global reference frame. The *state dynamics* — time derivatives of the state variables — encode the dynamic behavior of the aircraft in that global frame. A flight simulator integrates the state dynamics to generate a *state trajectory*: a record of how the state of the aircraft evolved over time.

This chapter develops state dynamics models for the paraglider system models. For each system model, it chooses a global coordinate system, defines a set of *state variables* $\boldsymbol{x}$ in terms of that global coordinate system, and defines the *state dynamics* $\dot{\boldsymbol{x}}$ in terms of the system dynamics.

7.1 State variables

To track the position of the glider, the state models must choose a reference point in the glider's local coordinate system. It does not have to be the same *reference point* used to calculate the system dynamics, but it turns out the riser midpoint RM is also good choice for tracking the glider position. Because the riser midpoint is close to where a pilot would likely mount their flight recorder, it is likely to be representative of the data in a flight track, which makes it the most convenient point for comparing real flight data to simulated data. Another advantage is that the riser midpoint is typically very close to the glider center of mass, which makes the position data easier to understand when developing the models.

Next, the state model must choose a coordinate system for the position. Most GPS applications, including paraglider flight records (IGC files), encode position using the WGS-84 *geodetic datum*, which uses the geocentric coordinates of latitude, longitude, and altitude. However, positioning on the global spheroid is overkill for these simulations, so to avoid the complexity involved with angular coordinates the state models here use a *tangent-plane* (tp)

approximation ([12], p. 27) that records position as a linear displacement from an arbitrary origin.

For orientation, there are two common representations: *Euler angles* and *quaternions*. Euler angles have the advantage of being easier to understand, but they can experience an issue known as *Gimbal lock* which prevents their use in situations where the aircraft rotates to extreme angles. Although the limitations of the paraglider aerodynamics make it unlikely for the simulator to encounter situations in which the glider is facing straight up or straight down, quaternions provide peace of mind and a minor improvement in computational efficiency.

Given these choices, the state variables of the 6-DoF models are four vectors:

$$\begin{aligned} \boldsymbol{r}_{RM/O} &\quad \text{absolute position of the riser midpoint } RM \\ \boldsymbol{v}_{RM/e} &\quad \text{translational velocity of the riser midpoint } RM \\ \boldsymbol{q}_{b/tp} &\quad \text{orientation of the body to the tangent plane} \\ \boldsymbol{\omega}_{b/e} &\quad \text{angular velocity of the body} \end{aligned} \tag{7.1}$$

Similarly, the 9-DoF models use the same four vectors, plus an additional quaternion and angular acceleration vector for the payload:

$$\begin{aligned} \boldsymbol{r}_{RM/O} &\quad \text{absolute position of the riser midpoint } RM \\ \boldsymbol{v}_{RM/e} &\quad \text{translational velocity of the riser midpoint } RM \\ \boldsymbol{q}_{b/tp} &\quad \text{orientation of the body to the tangent plane} \\ \boldsymbol{q}_{p/tp} &\quad \text{orientation of the payload to the tangent plane} \\ \boldsymbol{\omega}_{b/e} &\quad \text{angular velocity of the body} \\ \boldsymbol{\omega}_{p/e} &\quad \text{angular velocity of the payload} \end{aligned} \tag{7.2}$$

7.2 State derivatives

Next, define the derivatives of the state variables in terms of the current state and the system derivatives. The derivative of state variable for position is straightforward since it uses the same reference point as dynamics. The only modification is that the derivatives calculated by the system dynamics models were taken in the body and payload reference frames, $\mathcal{F}_b$ and $\mathcal{F}_p$, but tracking the position and orientation of the aircraft relative to the tangent plane requires derivatives taken with respect to the inertial frame $\mathcal{F}_e$. To provide the simulator with the proper derivatives, the state dynamics models must use the *equation of Coriolis* ([12], Eq. 1.4-2) to calculate the derivative of velocity taken with respect to the inertial frame:

$$
\begin{aligned}
{}^{e}\dot{\boldsymbol{v}}^{tp}_{RM/e} &= \boldsymbol{C}_{tp/b} \cdot \left({}^{b}\dot{\boldsymbol{v}}^{b}_{RM/e} + \boldsymbol{\omega}^{b}_{b/e} \times \boldsymbol{v}^{b}_{RM/e} \right) \\
{}^{e}\dot{\boldsymbol{\omega}}^{b}_{b/e} &= {}^{b}\dot{\boldsymbol{\omega}}^{b}_{b/e} \\
{}^{e}\dot{\boldsymbol{\omega}}^{p}_{p/e} &= {}^{p}\dot{\boldsymbol{\omega}}^{p}_{p/e}
\end{aligned}
$$

For the orientation state variable, the time derivative of a quaternion $\boldsymbol{q}$ that is tracking the orientation of an object can be calculated using the object's angular velocity vector $\boldsymbol{\omega} = \{p, q, r\}$ in the coordinate system attached to that object ($\boldsymbol{\omega}^{b}_{b/e}$ for the body, or $\boldsymbol{\omega}^{p}_{p/e}$ for the payload) ([12], Eq. 1.8-15):

$$
\boldsymbol{\Omega} \stackrel{\text{def}}{=} \begin{bmatrix} 0 & -p & -q & -r \\ p & 0 & r & -q \\ q & -r & 0 & p \\ r & q & -p & 0 \end{bmatrix}
$$

$$
\dot{\boldsymbol{q}} = \frac{1}{2} \boldsymbol{\Omega} \cdot \boldsymbol{q}
$$

The complete set of state dynamics equation for the 6-DoF models in terms of the system derivatives (6.2) and state variables (7.1) are then:

$$\begin{aligned}
{}^{e}\dot{\boldsymbol{r}}^{tp}_{RM/O} &= \boldsymbol{v}^{tp}_{RM/e} \\
{}^{e}\dot{\boldsymbol{v}}^{tp}_{RM/e} &= \boldsymbol{C}_{tp/b} \cdot \left({}^{b}\dot{\boldsymbol{v}}^{b}_{RM/e} + \boldsymbol{\omega}^{b}_{b/e} \times \boldsymbol{v}^{b}_{RM/e} \right) \\
{}^{e}\dot{\boldsymbol{q}}_{b/tp} &= \frac{1}{2}\boldsymbol{\Omega}_{b/tp} \cdot \boldsymbol{q}_{b/tp} \\
{}^{e}\dot{\boldsymbol{\omega}}^{b}_{b/e} &= {}^{b}\dot{\boldsymbol{\omega}}_{b/e}
\end{aligned} \tag{7.3}$$

Similarly, the complete set of state dynamics equation for the 9-DoF models in terms of the system derivatives (6.3) and state variables (7.2):

$$\begin{aligned}
{}^{e}\dot{\boldsymbol{r}}^{tp}_{RM/O} &= \boldsymbol{v}^{tp}_{RM/e} \\
{}^{e}\dot{\boldsymbol{v}}^{tp}_{RM/e} &= \boldsymbol{C}_{tp/b} \cdot \left({}^{b}\dot{\boldsymbol{v}}^{b}_{RM/e} + \boldsymbol{\omega}^{b}_{b/e} \times \boldsymbol{v}^{b}_{RM/e} \right) \\
{}^{e}\dot{\boldsymbol{q}}_{b/tp} &= \frac{1}{2}\boldsymbol{\Omega}_{b/tp} \cdot \boldsymbol{q}_{b/tp} \\
{}^{e}\dot{\boldsymbol{q}}_{p/tp} &= \frac{1}{2}\boldsymbol{\Omega}_{p/tp} \cdot \boldsymbol{q}_{p/tp} \\
{}^{e}\dot{\boldsymbol{\omega}}^{b}_{b/e} &= {}^{b}\dot{\boldsymbol{\omega}}^{b}_{b/e} \\
{}^{e}\dot{\boldsymbol{\omega}}^{p}_{p/e} &= {}^{p}\dot{\boldsymbol{\omega}}^{p}_{p/e}
\end{aligned} \tag{7.4}$$

The state dynamics models in (7.3) and (7.4) are ready to be used with a suitable numerical integration method to generate the state trajectories. Due to the significant nonlinear behavior of the dynamics, the implementation for this project uses a standard 4th order Runge-Kutta method.

Chapter 8

DEMONSTRATION

The motivation for this project was a need for paraglider flight dynamics models for commercial paraglider wings. The goal of this project was to build those system models by creating parametric *component models* that augment the limited available specifications with assumptions of the unknown structure. This chapter demonstrates one possible workflow to estimate the parameters of those component models by combining publicly available technical specifications and photographs with knowledge of typical paraglider wing design.

The paraglider wing used in this example is a Niviuk Hook 3. With forgiving flight characteristics targeting advanced beginners, this wing is not intended for acrobatics, so the *limitations* of the *aerodynamics method* are not an issue when simulating the majority of flights produced by this wing.

Wing data for a commercial wing is typically limited to four sources:

1. Technical specifications and user manuals
2. Flight test data from certifications and reviews
3. Pictures and videos
4. Physical measurements

For this chapter, only the first three will be utilized. Although physical measurements are ideal, they are frequently difficult to obtain (especially for older wings). Instead, this demonstration is focused on showing that it is feasible to create an approximate wing model even if physical measurements are unavailable.

Fig. 8.1: Front-view of an inflated Niviuk Hook 3

8.1 Technical specifications

The following sections demonstrate how to estimate the parameters for a size 23 version of the wing. The same process is used (but not shown) to create models of the size 25 and 27 wings to *validate* the modeling choices and implementation.

The process begins with the primary technical data from the official technical specifications manual:

Table 8.1: Wing data

Property [unit]	Size 23	Size 25	Size 27
Flat area [m^2]	23	25	27
Flat span [m]	11.15	11.62	12.08
Flat aspect ratio	5.40	5.40	5.40
Projected area [m^2]	19.55	21.25	22.95
Projected span [m]	8.84	9.22	9.58
Projected aspect ratio	4.00	4.00	4.00
Root chord [m]	2.58	2.69	2.8
Tip chord [m]	0.52	0.54	0.56
Standard mean chord [m]	2.06	2.14	2.23
Number of cells	52	52	52
Total line length [m]	218	227	236
Central line length [m]	6.8	7.09	7.36
Accelerator line length [m]	0.15	0.15	0.15
Solid mass [kg]	4.9	5.3	5.5
In-flight weight range [kg]	65-85	80-100	95-115

Recall that a "paraglider wing" includes both the canopy and the suspension lines, so the technical data describes both components. It also includes the weight range that the wing can safely carry while retaining control authority, which will be used to define a suitable payload.

8.2 Canopy

The first component model of the paraglider system is for the *canopy*. The canopy model combines an (idealized) *Foil geometry* model with physical details to estimate the aerodynamics and inertial properties of the canopy. For the canopy model parameters, it's easiest to think of them in two groups:

1. Parameters for the *design curves* that define the variables (3.15) of the *foil geometry model*.

2. Parameters for the physical details (5.10)

8.2.1 Foil geometry

Layout The first part of specifying a foil geometry is to layout the scale, position, and orientation of its sections.

For a parafoil, it's easiest to start by describing the geometry of the flattened (uninflated) canopy before dealing with the arc. This approach is made much easier by the choice of the *Simplified model* to define the *section index* as the normalized distance along the yz-curve. When a parafoil is flattened the section index corresponds to the normalized distance along each semispan, which allows the x-positions and chord lengths to be measured directly without regard for the arc.

First, consider the chord length distribution $c(s)$. The technical specifications only list the root, tip, and mean chord lengths, so more information is required. Thankfully, for parafoils a reasonable guess is that the wing uses a truncated elliptical distribution. (Paragliding wings commonly use truncated elliptic functions because they encourage elliptical lift distributions, thus reducing induced drag.) Such a truncated elliptical distribution can be easily parametrized by the wing root and wing tip section chord lengths, as shown by the *Elliptical chord* design curve. The technical specs list these two parameters as $c_{\text{root}} = 2.58$ and $c_{\text{tip}} = 0.52$, respectively. Using those values produces a standard mean chord length of 2.06, which exactly matches the value listed in the manufacturers specs,

so the assumption was justified. An additional check is to compare the area of the flattened chord surface projected onto the xy-plane; for these values the truncated elliptical produces a flattened area of 22.986 compared to the true specification of 23.0, which further confirms the design. (The small discrepancy may be explained by differences in measuring methodology or by the current absence of any geometry twist, but in practice the effect is negligible.)

Next is the fore-aft positioning of the sections, which are controlled by the $r_x(s)$ and $x(s)$ design curves. Although traditional wing geometry models would effectively choose $r_x(s) = 0$ and measure the x-offsets of each section's leading edge, that choice often produces an unnecessarily complicated $x(s)$ function. Instead, paragliders can often be described with constant $r_x(s)$ and $x(s) = 0$. As with the chord lengths, the value of $r_x(s)$ is easiest to estimate from the flattened wing; in fact, flattened drawings are commonly available in technical manuals, making them especially convenient. (Admittedly, such drawings do not always maintain the true aspect ratio, and so should be used with caution.) For this wing, a small amount of trial and error using a top-down view from the wing user manual suggests a constant $r_x(s) = 0.7$ gives a strong agreement with the drawing in the manual, as seen in Fig. 8.2.

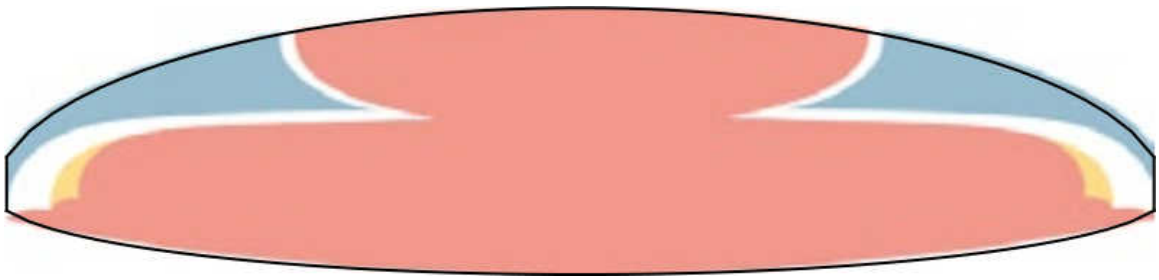

Fig. 8.2: Top-down outline of flattened canopy

The black outline is the boundary of the model's flattened chord surface. The colored background is taken from the user manual for the wing.

With the flattened chord surface completed, the next step is to define the *arc* (position in the yz-plane) to bend the flattened surface into its characteristic shape. Photos of the wing suggest that an elliptical arc segment is likely. The exact value of the arc reference points $r_{yz}(s)$ has a minimal impact for typical parafoils (which have relatively small geo-

metric twist), but a reasonable guess is to use the quarter-chord position $r_{yz}(s) = 0.25$. For the arc positions $yz(s)$, an *Elliptical arc* can be defined using three parameters: two for the normalized shape (Γ_{tip} and ϕ_{tip}) and one for the scale (b_{flat}). There are several ways to estimate the elliptical arc parameters of the physical wing, such as the width to height ratios, or visual estimation of the arc angle, but since the specs included both the flattened and projected spans, the simplest method is to guess a value for ϕ_{tip} and increase Γ_{tip} until the projected span matches the expected value. Starting with an initial guess of $\phi_{\text{tip}} = 75$, a few iterations shows good agreement with $\Gamma_{\text{tip}} = 32$. Checking the fit shows a projected span of $b = 8.845$ (versus the true value of $b = 8.840$) and a projected area of $S = 19.405$ (versus the true value of $S = 19.550$). As with the flattened values, the small discrepancy may be explained by differences in measurement methodology, and likely isn't worth optimizing further.

After the relatively straightforward process of positioning the sections is the more difficult task of estimating their orientation. In the *simplified model*, section roll $\phi(s)$ is defined by the curvature of the yz-curve and the section yaw $\gamma(s)$ is defined as zero, but the section pitch $\theta(s)$ (or *geometric torsion*) can be difficult to determine (even with a physical wing in hand). Relying on the fact that parafoils commonly benefit from a small amount of increasing geometric torsion towards the wing tips (or *washin*), a conservative guess of 4° at the wingtip should be reasonably accurate [32]. For lack of better information, this demonstration chose a piecewise *linear model* that grows 0–4° degrees over the range $0.05 \leq |s| \leq 1$.

Profiles Having finished defining the section layout (scale, position, and orientation), each section must be assigned an airfoil [37]. The most accurate way to determine the section profiles would be to cut open the wing and trace the outline of the internal ribs, but in this case that's not an option. Another option would be to search an airfoil database, but the simplest approach is to use a choice from literature. When using literature, it's important to keep in mind that although papers discussing "parafoils" and "ram-air parachutes" have much in common with paraglider canopies, those papers are typically analyzing large

canopies designed for heavy payloads.

From the ram-air category, [29] observes that many "older designs" use a Clark-Y airfoil with 18% thickness; it also mentions that "newer gliders" have been design with "low-speed sections", such as the LS(1)-0417 (for example, see [45]). For literature targeting paragliders specifically, one option is the NACA 23015: a classic, general purpose airfoil used in the *wind tunnel model* [19]. Another paraglider-specific option is the "Ascender": an 18% thickness airfoil developed for an open-design paraglider [32]; for an example of literature using that airfoil, see [46].

The criteria for selecting an airfoil is beyond the scope of this demonstration, but a key observation is the tendency for paragliders to use unusually thick airfoils. The reason for this is that thick airfoils tend to have more gentle stall characteristics, since their low-curvature leading edges encourage flow attachment as the angle of attack increases. Higher performance wings may select thinner airfoils to reduce drag, because the Hook 3 is a beginner-friendly wing this model uses a NACA 24018; it's similar to the 23015 used by the wind tunnel model but with 18% thickness. (For the curious reader, using the Ascender airfoil barely changes the equilibrium conditions for the wing; small changes to the equilibrium pitch angles and a small increase in the range of airspeeds, but otherwise the change had a surprisingly small effect.)

After choosing an airfoil, the next step is to modify it support the brake inputs. The unmodified airfoil defines the section profiles when no brakes are applied, but a paraglider must deform those profiles in order to turn and slow down. This poses a significant difficulty with modeling a paraglider, since the deformation is a complex process. Unlike wings made from rigid materials with fixed-hinge flaps, the brakes produce a continuous deformation along variable-length sections of the profile. Instead of dealing with that complexity, this project uses a strategy to simply guess the deflected geometry.

To begin, observe that the trailing edge of a braking paraglider typically exhibits a transition region followed by a gentle curve. In the interest of practicality, model the transition and trailing regions as circular arc segments. (This modeling choice is made with no theoretical justification beyond the recognition that spherical shapes tend to appear

as the energy-minimizing state of a flexible surface under tension.) Because this is not a theoretically well-justified model the algorithm will not be covered in detail, but this "two-circle model" can be used to generate a set of deflected airfoils.

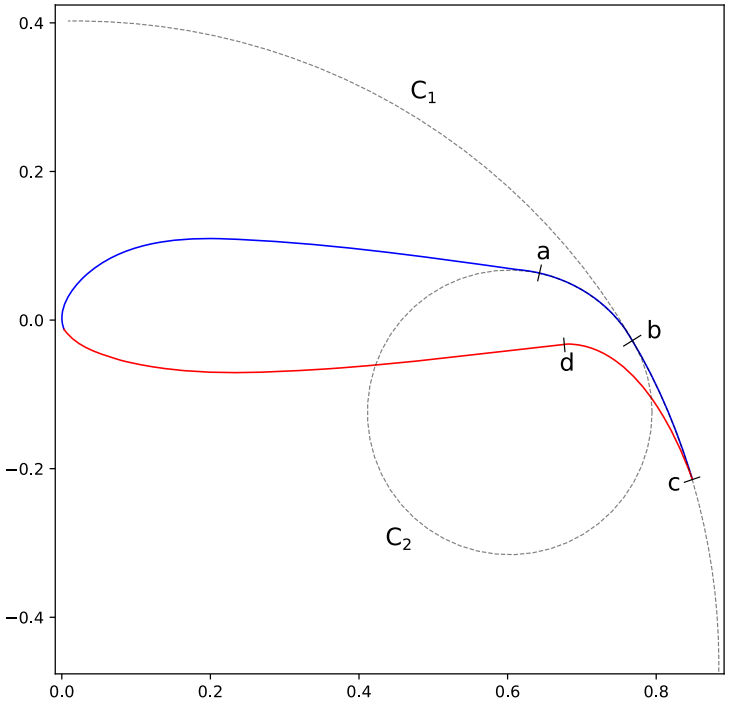

Fig. 8.3: Two-circle model to generate an airfoil with a smoothly-deflecting trailing edge.

For the upper surface, first choose a point (`a`) at some distance from the trailing edge (`c`) and attach a circle `C2` tangent to the airfoil at `a` and replace the transition region of the airfoil with an arc from `a` to `b`; then, place a second, larger, circle `C1` tangent at `b` and draw another arc for the remaining length of the upper curve. For the lower surface, choose a point `d` some distance roughly equal to the modified length of the upper surface and use a Bézier curve to draw a deflected lower surface between `d`, the new trailing edge `c`, and the point where the deformed upper surface curve crosses the original (undeformed) lower surface curve. The radius of the smaller circle `C2` controls the sharpness of the transition, and the radius of the larger circle `C1` controls the maximum steepness at the trailing edge. This procedure maintains the length of the upper surface, but neglects the wrinkling that normally occurs along the lower surface.

Using this procedure with the NACA 24018 as the baseline produces a set of reasonable-looking curves:

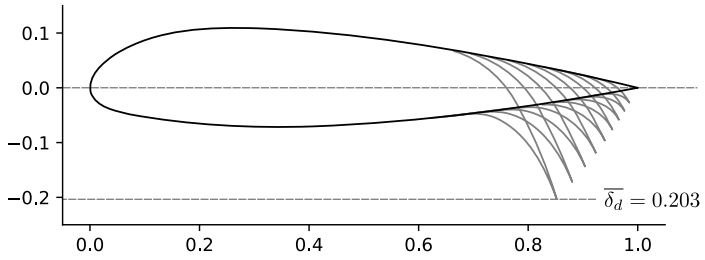

Fig. 8.4: Set of NACA 24018 airfoils with trailing edge deflections.

At this point the reader should be highly skeptical of this airfoil set. The choice of airfoil, and how the airfoil deforms in response to trailing edge deflections, is full of assumptions. Nevertheless, these results will be used for the remainder of this chapter as a means to demonstrate the working of the model. As a result, an important thing to keep in mind when interpreting the results of these choices is that choosing such a large radius for `C2` is wildly optimistic, but was chosen anyway to reduce the curvature of the transition region. For small brake inputs the transition curvature is negligible, but becomes progressively sharper as deflection increases. High curvature can be a problem for some theoretical models used to estimate the section coefficients (including the viscous/inviscid coupling method in XFOIL [47]), since the high curvature inhibits the method from converging on a solution when viscosity is taken into account. Softening the curvature allows the estimate to converge, but at the cost of hiding convergence failures that typically suggest flow separation. As a result, this profile set is likely to overestimate lift and underestimate drag.

8.2.2 Physical details

In addition to a *foil geometry*, a canopy model requires details of physical attributes such as surface material densities and air intake extents in order to calculate inertial properties and viscous drag corrections.

Surface materials In this case, the surface material densities can be read directly from the materials section of the user manual:

Table 8.2: Hook 3 material densities

Surface	Material	Density $\left[\frac{kg}{m^2}\right]$
Upper	Porcher 9017 E77A	0.039
Lower	Dominico N20DMF	0.035
Internal ribs	Porcher 9017 E29	0.041

In addition to the material densities, the canopy model requires the number of cells to determine the distribution of mass for the internal ribs. The specs lists $N_{\text{cells}} = 52$, which implies the wing has 53 ribs (including the wing tips). In reality the ribs are *ported* (holes that allow air to flow between cells) so assuming solid ribs is an overestimate, but since the canopy model is neglecting the mass from the remainder of the internal structure the discrepancy should (partially) balance out.

For the air intakes, the model must know the spanwise extent (since sections near the wing tips typically do not include air intakes). The user manual provides a projected diagram (Fig. 11.4, p. 17) which shows that the air intakes start at the 21st of 26 ribs (the 27th "rib" in the diagram is part of the stabilizer panel) spreading out from the central rib; assuming a linear spacing of the ribs this would correspond to $s = 0.807$, so $s_{\text{end}} = 0.8$ is a reasonable guess.

The other dimension of the air intakes is the size of their opening, which is determined by the extent of the upper and lower surface for each section profile. This value is difficult to determine precisely from photos, but thankfully its effect on the solid mass inertia and viscous drag is relatively minor; in the absence of physical measurements, a reasonable guess is $r_{\text{upper}} = -0.04$ and $r_{\text{lower}} = -0.09$ for an air intake length roughly 5% of the length of the chord. For a related discussion, see [46].

At this point the canopy can compute the total mass, which is another opportunity to sanity check the approximations. The technical specs list the total wing weight at 4.9kg,

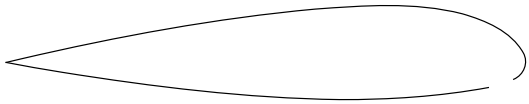

Fig. 8.5: NACA 24018 with air intakes

but the canopy materials included in this model only account for 2.95kg. This highlights the fact that the model neglects the extra mass due to things like the lines, riser straps, carabiners, internal v-ribs, horizontal straps, tension rods, etc. Fortunately, a significant amount of that missing mass is near the system center of mass and does not impart a major weight moment, so for the goals of this project the discrepancy is assumed to have a negligible impact on the overall system behavior.

Viscous drag corrections The last step is to add the empirical corrections to the section viscous drag coefficients. The first is a general factor applied to all the sections evenly to account for "surface characteristics", as estimated during wind tunnel measurements of parafoils in [43]:

$$C_{d,\text{surface}} = 0.004 \tag{8.1}$$

The second correction is to account for the additional viscous drag due to the presence of air intakes at the leading edge of some of the sections. In [30] they propose a simple linear relationship between the length of the air intake:

$$C_{d,\text{intakes}} = 0.07\,\frac{h}{c} \tag{8.2}$$

where h is the length of the air intakes and c is the length of the chord. This model assumes the air intakes constant (but proportional) size along the entire span between from $-s_{\text{start}} \leq s \leq -s_{\text{start}}$. As seen in Fig. 8.5, the air intakes are roughly 5% of the chord, for a value of roughly $C_{D,\text{intakes}} = 0.0035$. (The precise value is computed automatically by the implementation.)

8.3 Suspension lines

The second component model of the paraglider system is for the *suspension lines*. The behavior of the lines is deceptively complex, so the numerous parameters of the model were grouped by related functionality to (hopefully) make their relationships more intuitive.

8.3.1 Riser position

The first group of parameters (5.24) for the suspension line model determine the position of the harness (and pilot) underneath the canopy as a function of δ_a, the control input for the *Accelerator*.

Typically the most straightforward parameter to procure is κ_z: the vertical distance from the riser midpoint to the canopy as a ratio of the central chord c_{root}; for this wing, the technical specs listed this value as the "Central line length" and can be used directly, so $\kappa_z = \frac{6.8\,[m]}{2.58\,[m]} = 2.64$. Similarly, the accelerator line length (the maximum amount the accelerator can decrease the length of the central A lines) can also be read directly from the technical specs as $\kappa_a = 0.15\,[m]$.

Next, the canopy connection positions of the A and C lines as fractions of the central chord, κ_A and κ_C, are frequently visible in the line diagrams of the user manual; a quick measurement of the "Line plan" diagram (Sec. 11.4, p. 17) suggests $\kappa_A = 0.11$ and $\kappa_C = 0.59$.

The remaining parameter, κ_x, determines the fore-aft position of the riser midpoint. At first glance, this value can seem elusive, since it is difficult to determine precisely using any of the data in the technical manual; in fact, this value is also difficult to measure accurately from the physical wing, diagrams, or pictures. However, a useful strategy is to simply delay fixing the value of this parameter until the glider model is complete. The key insight is to recognize how the position of the harness impacts the equilibrium pitch angle of the wing, which in turn affects the equilibrium glide ratio of the complete glider. A simple rule of thumb is that modern paragliders are designed to maximize their glide ratio at "trim" conditions (that is, when no controls are being used), so choosing a value for κ_x

can be accomplished iteratively by choosing the value that maximizes the glide ratio with zero control inputs. If maximum glide requires braking, increase $kappa_x$; if maximum glide requires accelerating, decrease $kappa_x$. The exact value will depend on the type of harness and the weight limit the designer was using as the optimization target, but a reasonable starting point is $\kappa_x = 0.5$.

8.3.2 Brakes

The second group of parameters (5.25) for the suspension line model determine how the trailing edge of the canopy is deflected as a function of $\{\delta_{bl}, \delta_{br}\}$, the control inputs for the *Brakes*.

The first four parameters determine how the deflection distribution develops along the trailing edge as the brake lines are pulled. (Recall that the brake distribution is centered about s_{start} and s_{stop}, which are interpolated between their zero- and maximum-brake values.) Estimating these parameters starts by finding a view of the trailing edge when brakes are being applied:

First, the zero-brake values. From this picture the deflection appears to begin near the middle of each semispan. Adding a symmetric margin softens the distribution while keeping the starting point centered at $s = 0.5$, so $s_{\text{start},0} = 0.3$ and $s_{\text{stop},0} = 0.7$ look about right.

The maximum-brake values are more difficult, since they must coordinate with the value of κ_b, but from safety training footage it can be seen that maximum brakes produce a deflection from roughly $s_{\text{start},1} = 0.08$ to $s_{\text{stop},1} = 1.05$ (where the stopping position exceeds the wing tip to indicate that the wing tip itself experiences a small deflection).

Next, the model needs the maximum distance the brake lines can be pulled. On a real wing the brake lines effectively don't have a well-defined limit, since a pilot can literally wrap the brake lines around their hand to pull the trailing edge all the way back to the risers, but in practice the airfoil set Fig. 8.4 that defines the deflected profiles is limited to some maximum deflection distance. For that reason, the *Suspension lines* model uses brake inputs on a scale from 0 to 1, with a maximum brake deflection distance κ_b. The value of κ_b

Fig. 8.6: Rear-view of an inflated Hook 3 with symmetric brake deflections

should maximize the usable range of the brakes without causing the normalized deflection distance $\bar{\delta}_d$ (5.1) of any section to exceed the distance supported by the airfoil set. Written as an optimization in terms of (5.14), the goal is to calculate the value of κ_b such that:

$$\max_s \frac{\delta_d(s,1,1)}{c(s)} = \bar{\delta}_{d,\max}$$

Checking the airfoil set used for this model (Fig. 8.4), define $\bar{\delta}_{d,\max} = 0.203$. Solving the optimization problem determines $\kappa_b = 0.426\,[m]$. This procedure is unfortunately convoluted, but in summary: for this specific airfoil set, the foil's chord distribution, and these brake position parameters, the model can allow the brake lines to be pulled a maximum distance of $42.6\,[cm]$.

To check the model fit, plot the undeflected and deflected trailing edge to compare with the reference photos:

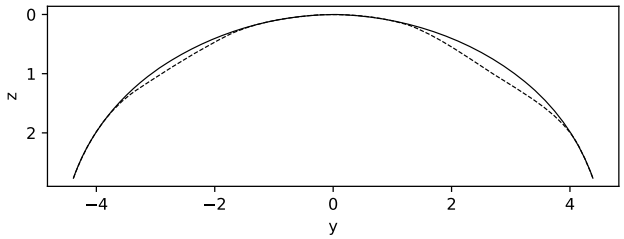

Fig. 8.7: Niviuk Hook 3 23 brake distribution, $\delta_{bl} = 0.25$ and $\delta_{br} = 0.5$

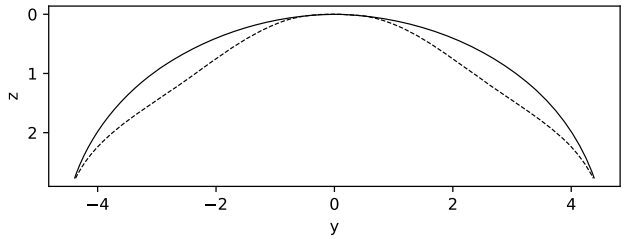

Fig. 8.8: Niviuk Hook 3 23 brake distribution, $\delta_{bl} = 1.00$ and $\delta_{br} = 1.0$

8.3.3 Line drag

The third group of parameters (5.26) for the suspension line model determine the aerodynamic drag of the lines. Because the model is focused on providing functionality instead of a detailed (and tedious) layout of every line, it computes the drag by lumping the total area of the lines into a small number of points. For this demonstration, satisfactory results can be achieved with just two points (one for each semispan) and crude estimates of the true line area distribution.

First, the total line length for this wing is listed directly in the technical specs, $\kappa_L = 218\,[m]$. Next, κ_L must be multiplied by the average diameter of the lines κ_d to get their total surface area. Although a complete set of diameters for each line segment are given in the "Lines Technical Data" section, computing an accurate distribution would require their detailed layout; instead, with lower sections of the cascade averaging $2.8\,[mm]$ and upper sections using $0.6\,[mm]$ lines, a good starting point is to assume an average diameter of $\kappa_d = 1\,[mm]$. Next, the area is divided into the two control points, which must be positioned at the area centroids of their group of lines. For an approximate model such as

this, the positions of the points are easiest to estimate visually; using Fig. 8.6 they appear to be around $\boldsymbol{r}_{CP/R} = \langle -0.5c_{\text{root}}, \pm 1.75, 1.75 \rangle$. Lastly, each lumped line area is assigned a drag coefficient; because the lines are essentially cylinders, a suitable drag coefficient is simply $C_{d,l} = 1$ [20].

8.4 Payload

The final component model of the paraglider system is for the *harness*. This component is responsible for positioning the mass of the payload (harness and pilot) as a function of weight-shift, and computing the aerodynamic drag applied to the payload. The parameters of the model are the total mass of the payload (m_p), the vertical distance of the mass centroid below the riser midpoint (z_{riser}), the cross-sectional area of the payload (S_{payload}), the aerodynamic drag coefficient ($C_{d,\text{payload}}$), and the maximum horizontal distance a pilot can displace the centroid using weight-shift control (κ_w).

For the total mass, the technical specs list the weight range for the size 23 wing as 65–85 [kg], so $m_p = 75\,[\text{kg}]$ is a conservative choice.

For the mass centroid, one option is to consider the DHV airworthiness guidelines [48], which specify that the riser attachment points must be "35–65cm above the seat board", which suggests that $z_{\text{riser}} = 0.5\,[\text{m}]$ is a reasonable value in most cases. Alternatively, simply look up the technical diagram of a suitable harness; for example, the wing certification flight tests (published in the Hook 3 User Manual, p. 22) list the "harness to risers distance" as 49cm.

For the surface area and its associated drag coefficient, consider [31] (p. 85) or [30] (p. 422); for a 75kg payload with a non-pod-style harness, a reasonable estimate of the area would be $S_{\text{payload}} = 0.55\,[\text{m}^2]$ with an drag coefficient of $C_{d,\text{payload}} = 0.8$.

Lastly, when choosing a weight shift limit, underestimates are preferable to overestimates, since an underestimate merely limits the range of behavior the model can produce, whereas an overestimate can produce fictitious behavior; in the absence of a rigorous measurement, a conservative guess is $\kappa_w = 0.15\,[\text{m}]$.

Chapter 9

VALIDATION

9.1 Foil aerodynamics

The *Foil aerodynamics* chapter selected Phillips' NLLT because it appeared to satisfy the *Modeling requirements* established at the beginning of this paper; this section uses wind tunnel measurements to validate that choice. First it recreates the geometry using the *Simplified model*, then it recreates the range of test conditions used by the experiment and tabulates the aerodynamic coefficients estimated by the NLLT. The estimates are compared to the wind tunnel data, as well as to other standard aerodynamic models commonly recommended for nonlinear geometries.

9.1.1 Geometry

The geometry from a 2015 parafoil wind tunnel test [19] makes an excellent case study of a foil specification from literature that positions the sections using alternative reference points on the section chords. Moreover, the geometry satisfies the assumptions of the *Simplified model*, making an implementation of the geometry almost trivial.

First, the paper describes the geometry of the full-scale canopy they wish to study:

Table 9.1: Full-scale wing dimensions

Property	Value	Unit
Arch height	3.00	m
Central chord	2.80	m
Projected area	25.08	m^2
Projected span	11.00	m
Projected aspect ratio	4.82	–
Flat area	28.56	m^2
Flat span	13.64	m
Flat aspect ratio	6.52	–

For the wind tunnel test, a one-eighth scale physical model was constructed from a wood-carbon frame with polyurethane foam sections covered in fiberglass. Physical dimensions and positions were provided for the physical model as pointwise data with linear interpolation between each point.

Table 9.2: Wind tunnel wing geometry data at panel's ends

i	y [m]	z [m]	c [m]	r_x	r_{yz}	θ [deg]
0	-0.688	0.000	0.107	0.6	0.6	3
1	-0.664	-0.097	0.137	0.6	0.6	3
2	-0.595	-0.188	0.198	0.6	0.6	0
3	-0.486	-0.265	0.259	0.6	0.6	0
4	-0.344	-0.325	0.308	0.6	0.6	0
5	-0.178	-0.362	0.339	0.6	0.6	0
6	0.000	-0.375	0.350	0.6	0.6	0
7	0.178	-0.362	0.339	0.6	0.6	0
8	0.344	-0.325	0.308	0.6	0.6	0
9	0.486	-0.265	0.259	0.6	0.6	0
10	0.595	-0.188	0.198	0.6	0.6	0
11	0.664	-0.097	0.137	0.6	0.6	3
12	0.688	0.000	0.107	0.6	0.6	3

It is important to notice the difference between the section numbers i used in the paper and the section indices s used in the simplified model; the section indices are easily calculated using the normalized linear distance along the $\langle y, z \rangle$ points. Also, the reference data is defined with the wing tips at $z = 0$, whereas the convention of this paper places the canopy origin at the leading edge of the central section; this is easily accommodated by subtracting the central $z = -0.375$ from all z-coordinates. (Alternatively, the `implementation` of the simplified model in `glidersim` can shift the origin automatically.)

Calculating the section indices for each point and using linear interpolation as a func-

Fig. 9.1: NACA 23015

tion of the section index produces a set of piecewise-linear design curves, and assigning every section a NACA 23015 airfoil (Fig. 9.1) completes the foil geometry model.

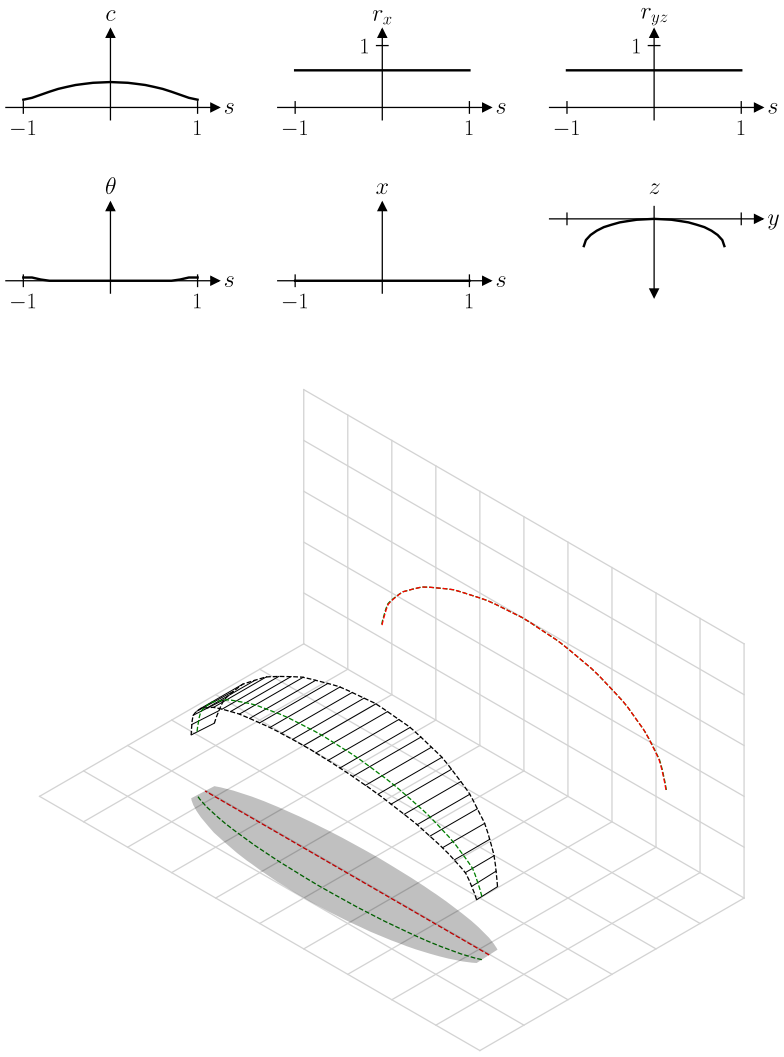

Fig. 9.2: Chord surface for Belloc's reference paraglider wing.

9.1.2 Wind tunnel setup

The setup mounted the 1/8-scale model on a 1 meter rod connected to force sensors, and set the wind tunnel to a 40 m/s airspeed. Measurements were taken with the angle of attack and sideslip ranging over $-5 < \alpha < 22$ and $-15 < \beta < 15$ (a range suitable capturing longitudinal performance post-stall). For better accuracy, wind tunnel measurements should be corrected for wall interactions with the flow ([49]; [8], Sec. 10.3). However, because classical wind tunnel wall corrections assume a flat wing, the data for the arched parafoil

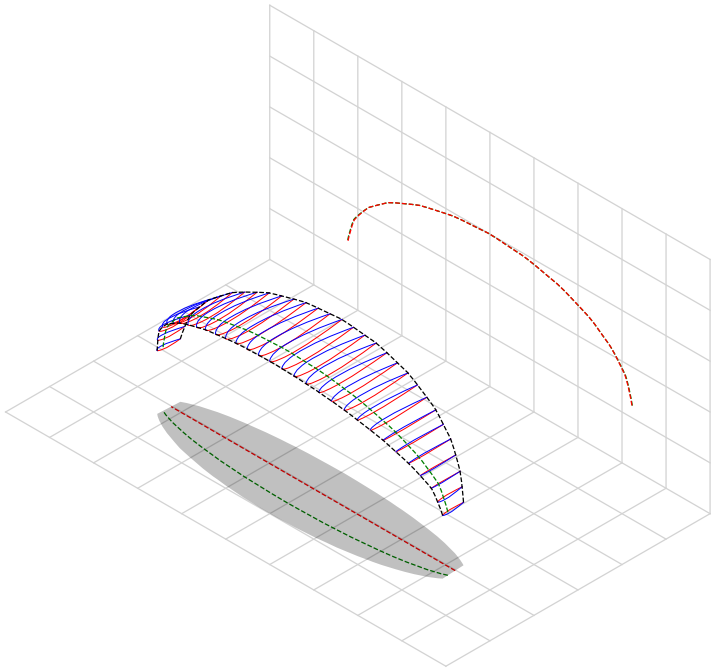

Fig. 9.3: Profile surface for Belloc's reference paraglider wing.

are uncorrected for wall effects.

9.1.3 Aerodynamics models

The wind tunnel data will be compared to three theoretical aerodynamics models, one that includes viscous effects, and two that do not (inviscid models):

1. NLLT: the *numerical lifting-line* model from [21]

2. AVL: an extended vortex lattice method by Mark Drela [41] (who also authored XFOIL [47] while at MIT) . With a long history in academic research, this is the primary reference for comparing the results of the NLLT.

3. XFLR5: an experimental vortex lattice method from the open source wing modeling tool by André Deperrois. This model is marked "experimental" by the author because it is still under development, but the principle is to mitigate the "small angles" approximation relied on by standard vortex lattice methods by reorienting the foil geometry instead of reorienting the flow. The purpose of including this method in

these tests is to show the effect of the simplifying assumptions used when designing the system of equations for aerodynamics models. For conventional aircraft where the flow angles are relatively small, small angle approximations are reasonable, but for nonlinear geometries at large angles of attack, classic methods such as AVL begin to struggle.

9.1.4 Results

9.1.4.1 Lift vs drag The standard way to summarize the efficiency of a wing is to plot the amount of lift it produces versus the amount of drag; with practice, such charts can be used to quickly approximate performance characteristics such as its glide ratio. They are also useful for quickly comparing the relative performance of each aerodynamics method.

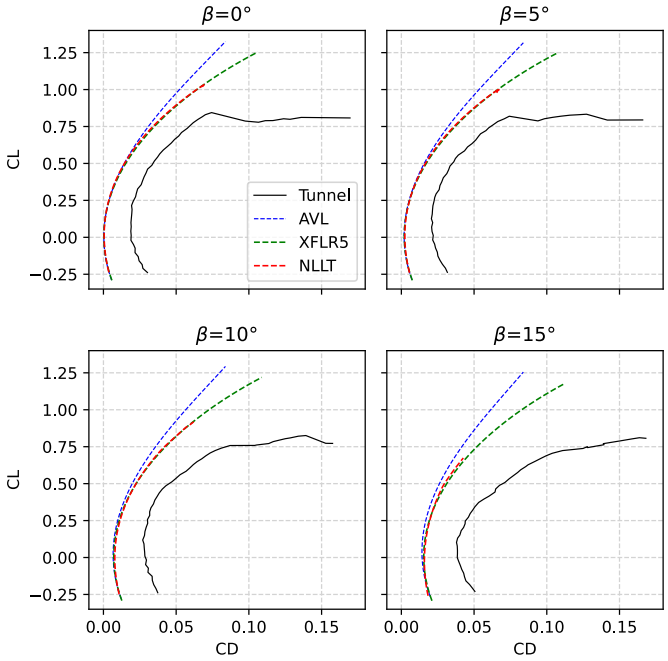

Fig. 9.4: Lift vs induced drag

The first thing step during validation is to verify the test setup for each of the models. One way to do that is by comparing methods that are expected to produce equivalent results; in this case, the inviscid methods from AVL and XFLR5 should be nearly identical

at low angles of attack, and should estimate zero drag at zero lift coefficient and zero sideslip. Because the NLLT uses aerodynamic coefficients that include viscous effects it is not directly comparable to the inviscid models, but because viscosity is not expected to have a significant effect on lift at low angles of attack, it is possible to disregard the viscous drag coefficients and plot the pseudo-inviscid polar curve by setting the viscous drag coefficients to zero, as shown in Fig. 9.4. (This is a "pseudo" inviscid curve since the section lift coefficients used by the NLLT include viscous effects.) The resulting drag coefficient is limited to drag produced by the creation of lift, as would be predicted by the inviscid methods. This plot is useful because it validates that the geometry model and test conditions were configured correctly in all tools, and provides evidence that the NLLT was implemented correctly.

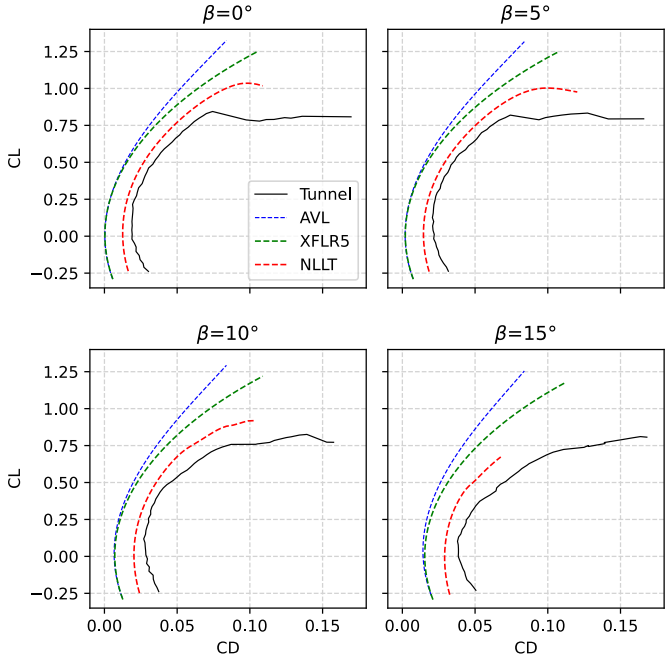

Fig. 9.5: Lift vs drag

The second plot (Fig. 9.5) compares the inviscid methods to the NLLT with the unadjusted aerodynamic coefficients from XFOIL. The first thing to note is the difference compared to the pseudo-inviscid plot (Fig. 9.4): as expected, including viscous drag sig-

nificantly improves the agreement between the theoretical and experimental results for the NLLT. Another observation is the significance of the inviscid assumption, with both inviscid methods overestimating lift and underestimating drag at higher angles of attack. This plot also appears to show the effect of the "small angles" approximation relied on by AVL, with the experimental "tilted geometry" method from XFLR5 providing better accuracy at high angles of attack and sideslip.

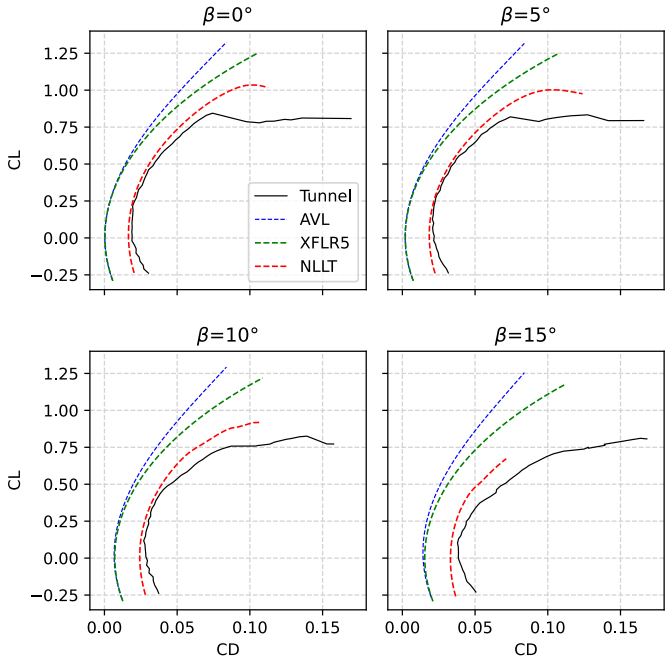

Fig. 9.6: Lift vs drag with extra viscous drag due to "surface characteristics"

A final plot (Fig. 9.6) is more for future reference than validation. Instead of the unadjusted aerodynamic coefficients from XFOIL, it adds the additional viscous drag due to "surface characteristics" suggested in [43] as a result of their wind tunnel tests on parafoils. Because this empirical adjustment will be used in the *Demonstration* portion of this paper, this plot is useful to show the expected accuracy of the NLLT when applied to a model of commercial paraglider wing used for dynamic simulations.

9.1.4.2 Coefficients vs angle of attack Another valuable way to summarize wing behavior is to plot the longitudinal-centric coefficients (lift, drag, and pitching moment) versus the angle of attack α. These results are grouped into four quadrants by the sideslip angle β used during the test.

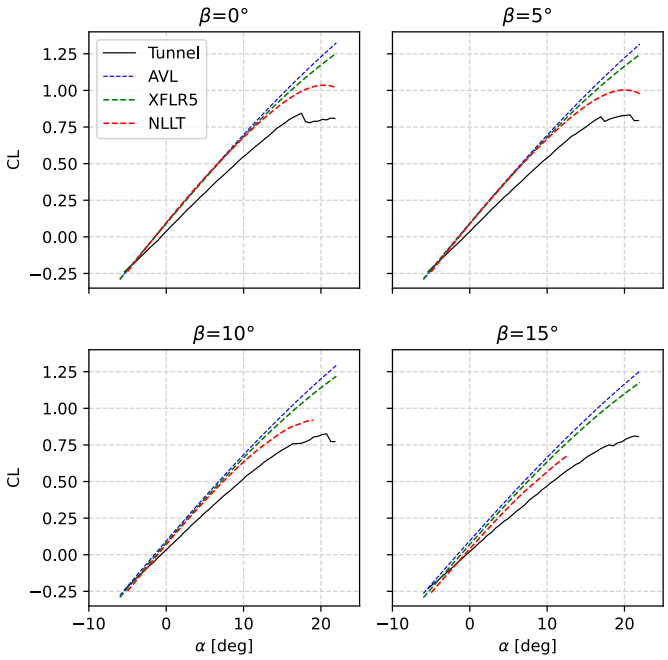

Fig. 9.7: Lift coefficient vs angle of attack

The first (and arguably most interesting) plot is for lift versus angle of attack (Fig. 9.7). Separating lift into its own plot reveals the source of the flatline region in the "Lift vs drag" plots; the wing enters stall (so lift ceases to grow) at approximately $\alpha = 17°, \beta = 0°$, and slightly earlier during sideslip (although the nonlinearity of the geometry dramatically affects the stall pattern and "smooths" the effect making it more difficult to see).

The more interesting result, however, is that all three theoretical methods are in very close agreement for the majority of the range, they all mispredict the zero-lift angle of attack, and they all uniformly overestimate the slope of the lift curve. This anomaly is difficult to explain; at $\beta = 0°$ and low angles of attack, the effects of viscosity should have a negligible effect on lift, and the vortex lattice methods should perform very well, but they

don't. The fact that the NLLT agrees with them is encouraging (again, the fact that it uses lift coefficients that account for viscosity should have a negligible effect in this test, and so the NLLT is expected to agree with the inviscid methods). I contacted the authors of both the wind tunnel data and the NLLT, and neither author had any immediate feedback on what would cause this issue. Nevertheless, there are two useful takeaways:

1. The NLLT is at least as accurate as the inviscid methods.

2. The NLLT is approximating the nonlinear effects of early stall, whereas the inviscid methods maintain a virtually linear response. This is an encouraging sign that the NLLT is a suitable choice given my *Modeling requirements* that the aerodynamics should provide "graceful degradation of accuracy" as it approaches high angles of attack.

This plot also highlights a limitation of relying on aerodynamic coefficients: the NLLT cannot produce a solution if any of the sections experience a section-local angle of attack that exceeds the range supported by the set of aerodynamic coefficients. This is effect is clear as the sideslip angle increases: because the wing is arched, as sideslip becomes positive (so the relative wind approaches from the right of the wing) the angle of attack on the left wingtip increases. As a result, as soon as global α and β produce a section-local α that exceeds the maximum value in the coefficients lookup table, the NLLT cannot produce a solution. The inviscid models, on the other hand, are founded on linear relationships with no upper bound, allowing them to generate estimates at significantly higher angles of attack and sideslip. Whether a bad estimate is better than no estimate, however, depends on the application.

When considering drag versus angle of attack (Fig. 9.8), the most noteworthy details are how all three methods fail to predict the rapid increase in drag as the wing enters the stall region, and how the "tilted geometry" of the XFLR5 model allows it to more accurately track the shape (if not the value) of the viscous solution.

Another coefficient that has a strong impact on the pitch stability of a paraglider canopy is the pitching moment versus angle of attack (Fig. 9.9). This plot can be viewed

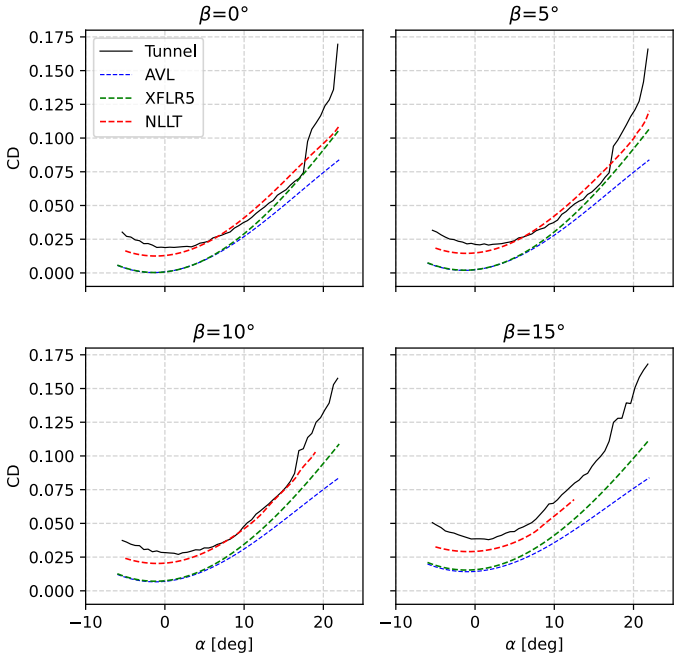

Fig. 9.8: Drag coefficient vs angle of attack

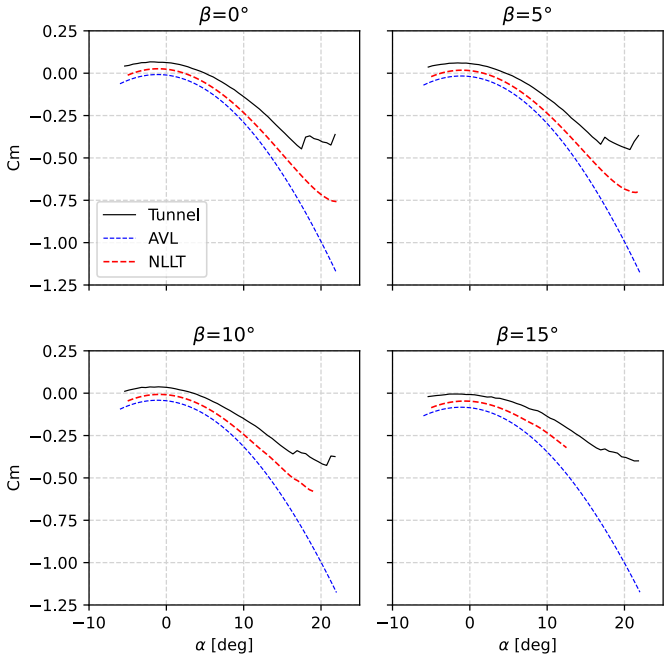

Fig. 9.9: Pitching coefficient vs angle of attack.

as pre- and post-stall conditions (before and after $\alpha = 17°$ in the $\beta = 0°$ quadrant), and are worth considering separately.

In the pre-stall region, the plot shows how a negative pitching moment grows with α, resulting in negative feedback that provides a restoring force back to equilibrium. If the wing pitches backwards, the negative pitching moment will help bring the canopy back overhead into a stable position.

In the post-stall region, the effect of flow separation can be seen in the experimental data by the sudden flat response of the pitching coefficient to α. This reason is complex, but informative:

- Because the lift vector at positive α points forwards, lift creates a negative (forward) pitching moment. At stall, lift decreases, which increases C_m.
- Because drag points backwards, it creates a positive (backwards) pitching moment. At stall, drag dramatically increases, which also increases C_m.
- At stall, flow separation typically starts at the trailing edge on the upper surface. The loss of pressure creates a negative (forwards) pitching moment, which decreases C_m.

For the wind tunnel model, it appears that (again, for the $\beta = 0°$ case) these effects are counteracting each other, producing a relatively flat C_m in the post-stall region. The inviscid method used by AVL fails to capture the nonlinearity of flow separation, causing it to overestimate the lift and underestimate drag that together producing a significantly inaccurate pitching moment post-stall. (Unfortunately the experimental method in XFLR5 had a bug that produced zero sideforce, so its results are omitted.) The NLLT performs much better, but still highlights the effect of using the well-known "optimistic" estimates produced by XFOIL near the stall region; and again, the NLLT fails to converge when the section-local α of the downwind wingtip exceeds the maximum α supported by the coefficients lookup table instead of producing progressively more incorrect results.

9.1.4.3 Coefficients vs sideslip A third perspective of wing behavior is to plot the coefficients that affect motion in the y-direction (sideforce, rolling moment, and yawing moment) versus angle of sideslip β. These results are grouped into four quadrants by the angle of attack α used during the test. Unfortunately, the experimental method in XFLR5 had a bug that produced zero sideforce, which is also coupled to the roll and yaw moments, so its results are omitted.

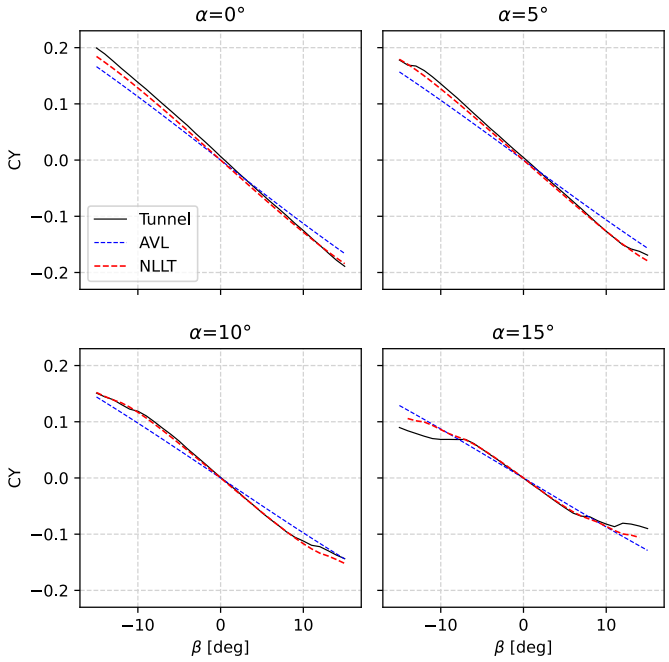

Fig. 9.10: Lateral force coefficient vs sideslip

Plotting sideforce vs sideslip (Fig. 9.10) showed good agreement between the experimental data and both theoretical models, although the NLLT has a slight accuracy advantage over the inviscid method.

In the rolling moment versus sideslip test (Fig. 9.11) we find the only examples where the inviscid method outperforms the NLLT, but otherwise this plot demonstrates no noteworthy effects.

The last plot, for the yawing moment versus sideslip (Fig. 9.12) has several similarities to Fig. 9.9, except instead of demonstrating the pitch stability of the wing, it demonstrates

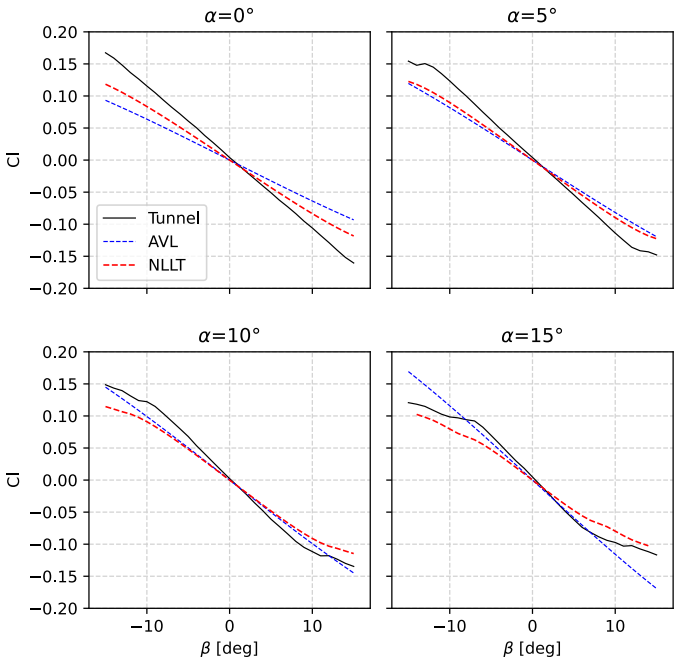

Fig. 9.11: Rolling coefficient vs sideslip

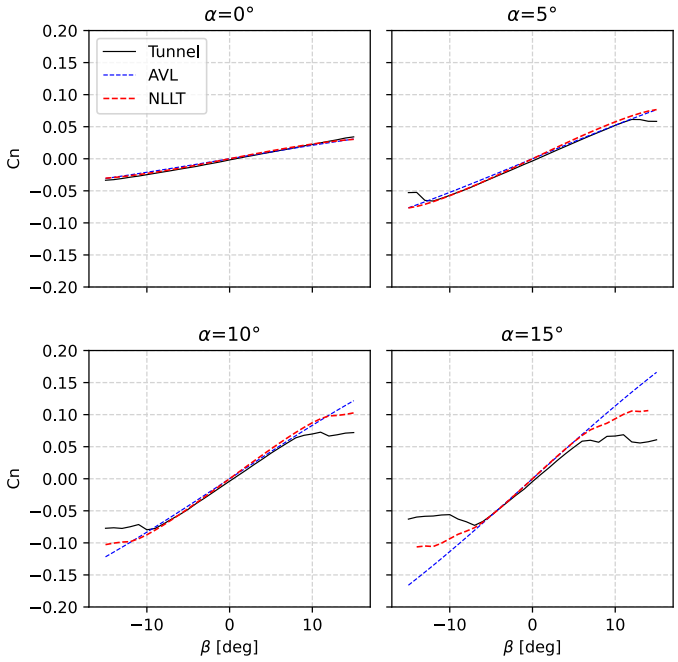

Fig. 9.12: Yawing coefficient vs sideslip

the yaw stability of the wing. When the relative wind approaches from the right ($\beta > 0°$) a positive yaw moment will turn the canopy into the wind, and vice-versa for wind from the left. And again, the effect of failing to accurately model stall conditions on individual sections (the downwind sections, specifically) causes both methods to overestimate the restoring moment. Nevertheless, the NLLT succeeded in capturing at least part of the effect, once again proving the value of the method over purely inviscid solutions.

9.2 Niviuk Hook 3 system dynamics

The previous chapter provided a *Demonstration* of how to estimate the parameters of the component models for a commercial paraglider wing. Having defined the component models, they are combined into a composite *System dynamics* model that provides the behavior of the complete glider. Getting to this point with such little information required many modeling assumptions, simplifications, approximations, and outright guesswork, so the natural next step is to question the validity of the model: how accurately does it estimate the true behavior of the physical system? In any modeling project it is vital to validate the model by comparing its estimates to experimental data, and this case is no exception.

Unfortunately, experimental data is extremely scarce for commercial paraglider wings. Unlike the previous section, wind tunnel measurements are unavailable. What's worse, the dynamic behavior of a wing in motion is significantly more complex than the static behavior of a wing held fixedly in a wind tunnel. As a result, validation is limited to point data and general expectations gleaned from sources such as glider certifications and consumer wing reviews. Clearly such sources lack the rigor to "prove" model accuracy, but — when taken together — they can still provide incremental confidence that a model is adequate to answer basic questions of wing performance.

9.2.1 Polar curve

The conventional way to summarize the performance of a gliding aircraft is with a chart called the *polar curve*. These curves show the vertical and horizontal speed of the aircraft at equilibrium over the range of brake and accelerator inputs, providing information such as the speed range of the glider and its glide ratio at different speeds. Given the wealth of information compactly communicated by a polar curve, they are an excellent starting point for critiquing the estimates of a flight dynamics model for a glider.

The previous section demonstrated the creation of a paraglider model for a Niviuk Hook 3, size 23. Now, models for the larger sizes of the wing (created using the same workflow) will be compared to experimental data by comparing measurements from test flights to the predicted polar curves.

9.2.1.1 Size 25 The experimental data for this section is taken from a size 25 version of the wing that was reviewed for the French magazine "Parapente Mag". Unfortunately, reviews such as this cannot provide the entire polar curve: because each point is laborious to measure accurately, reviews only provide noteworthy values, such as the minimum and maximum speeds, or the horizontal and vertical speeds that mark the "minimum sink" and "best glide" operating points of the glider. Despite this ambiguity, by plotting the experimental point data over the theoretical curve it is possible to get a sense of the general accuracy of the model estimates.

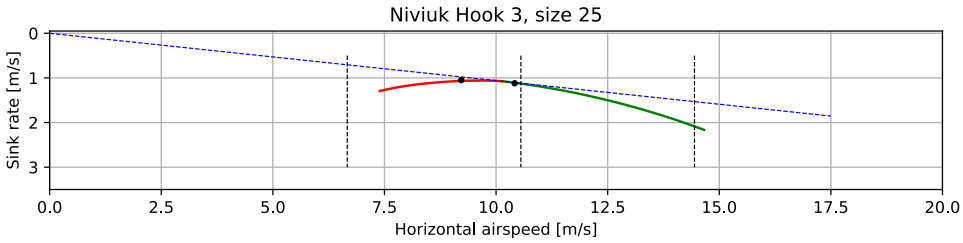

Fig. 9.13: Polar curve for Niviuk Hook 3 size 25

Colored markings are theoretical data from the model, black markings are experimental data from Parapente Mag. Red represents symmetric braking, green represents accelerating, and the blue diagonal line marks the predicted best glide ratio. The three black vertical lines mark the experimental values for minimum speed, trim speed, and maximum speed; the left black dot is the "minimum sink" operating point, and the right dot is the "best glide" operating point.

If the model is a good approximation of the glider that generated the data — and assuming the data was collected accurately — then the experimental values should match the predicted values:

- The minimum ground speed should align with the leftmost endpoint of the red curve
- Trim speed should align with the point where the red and green curves connect
- The maximum ground speed should align with the rightmost endpoint of the green curve
- The "minimum sink" operating point should lie on the point where the curve reaches its minimum
- The "best glide" operating point should lie on the point where the blue line touches the polar curve

Although the diagram is a convenient way to summarize so much information it can be hard to distinguish specific values, so their numerical equivalents are listed below.

Table 9.3: Niviuk Hook 3 25 simulated polar curve vs flight data

Value	Experimental	Simulated	Error
Minimum speed	6.7	7.4	+10%
Minimum sink <h, v>	9.22, 1.02	9.6, 1.06	+4.2%, +3.9%
Trim speed	10.6	10.2	-3.8%
Maximum speed	14.4	14.7	+2.08%
Best glide <h, v>	10.4, 1.12	10.2, 1.08	-1.9%, -3.6%
Best glide ratio	9.3	9.44	+1.5%

Observations:

- The minimum ground speed of the theoretical model is significantly higher than the experimental value. That may be explained by the conservative value of $\kappa_b = 0.44\,[m]$

(the maximum distance the brakes can be pulled; see the earlier discussion when defining the parameters for the *Brakes*). The review listed the maximum brake length as >60cm, which suggests that this model can only apply <73% of the full range of brakes, so this result in unsurprising.

- Minimum sink occurs at about 0.4 m/s slower ground speed. This may be related to the procedure to generate the deflected *Profiles*, to the deflection distribution, or to the aerodynamic coefficient estimates from XFOIL.
- Minimum sink rate is remarkably close (1.06 versus 1.02 m/s), which I find surprising since I expected the "optimistic" airfoil set Fig. 8.4 to overestimate lift during braking.
- The theoretical model underestimates the ground speed at trim. Although this could be due to it overestimating the drag, it is far more likely that the model is overestimating the lift of the wing, so less speed is required to counteract the weight of the glider.
- This experimental data reported the best glide at 10.4 m/s when trim was 10.6 m/s. This disagrees with our earlier assumption that best glide should occur at trim.
- The model overestimates the maximum ground speed. This may suggest it is underestimating drag, or it could suggest that the model parameters are wrong (κ_C in particular has a large impact on maximum speed), or it could be because this rigid body model neglects foil deformations (it assumes the accelerator produces a perfect pitch-rotation of the foil) as well as the section profile deformations that increase with speed.

In truth, these observations are just a few of the possible issues with the theoretical model (not to mention issues with the experimental data itself); there are so many simplifications at work, and point data cannot hope to reveal all their flaws. These results suggest that the performance of the model is excellent when predicting longitudinal equilibrium, but a wider variety of wing models need to be examined to determine if this excellence generalizes to other wings.

9.2.1.2 Size 27 The experimental data for this section is taken from a size 27 version of the wing that was reviewed for the Spanish magazine "Parapente". As with the size 25 model, plotting the experimental data on top of the theoretical curves produces valuable reference data:

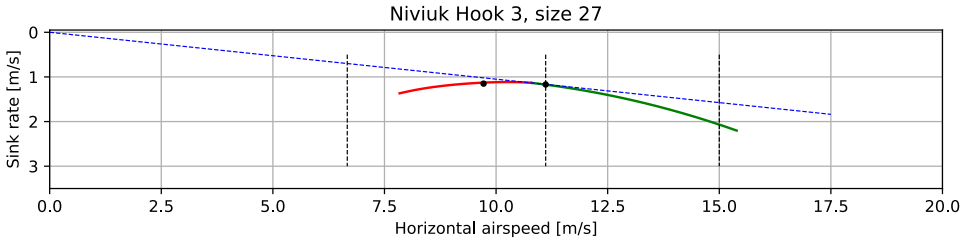

Fig. 9.14: Polar curve for Niviuk Hook 3 size 27

Colored markings are theoretical data from the model, black markings are experimental data from Parapente. Red represents symmetric braking, green represents accelerating, and the blue diagonal line marks the predicted best glide ratio. The three black vertical lines mark the experimental values for minimum speed, trim speed, and maximum speed; the left black dot is the "minimum sink" operating point, and the right dot is the "best glide" operating point.

As before, the numerical equivalents of the data in the figure above:

Table 9.4: Niviuk Hook 3 27 simulated polar curve vs flight data

Value	Experimental	Simulated	Error
Minimum groundspeed	6.7	7.83	+17%
Minimum sink <h, v>	9.72, 1.15	10.2, 1.12	+4.9%, -2.6%
Trim speed	11.1	10.8	-2.7%
Maximum speed	15	15.4	+2.7%
Best glide <h, v>	11.1, 1.17	10.8, 1.13	-2.7%, -3.4%
Best glide ratio	9.5	9.52	0.21%

The observations are similar to that for the size 25 model. Overall the fit is excellent. This model was limited to $\kappa_b = 0.46\,[m]$, or <76% of the usable ">60cm" brake length, so the minimum ground speed is still too high. And again, the model underestimates the ground speed at trim. The best glide ratio matches exactly, although the theoretical model

still slightly underestimates the ground speed where that occurs.

9.2.2 Pitch stability

Another simple sanity check is to verify the glider pitch stability by flying on a straight course at maximum speed and abruptly releasing the accelerator ([48], Sec. 4.1.5). Releasing the accelerator shifts the payload to shift aft, causing the canopy to pitch backwards; in the positive-pitch position the glider briefly ascends as it converts the energy from its high airspeed into altitude, but because the wing loses airspeed so quickly it will "overshoot" its equilibrium point and need to dive forward as the glider attempts to reestablish equilibrium.

The danger of this pitch-forward behavior is that it may induced a frontal collapse of the canopy. To estimate the safety margin of the wing, the test assigns a grade based on the negative pitch angle as it dives forward. If the wing pitches forward less than 30° it receives an "A"; if it pitches forward 30–60° it receives a "C", and for >60° it receives an "F". The Niviuk Hook 3 is rated as an "B" wing, and should not pitch forward more than 30°. Using this model to simulate the test protocol by releasing the accelerator in 0.3s produces:

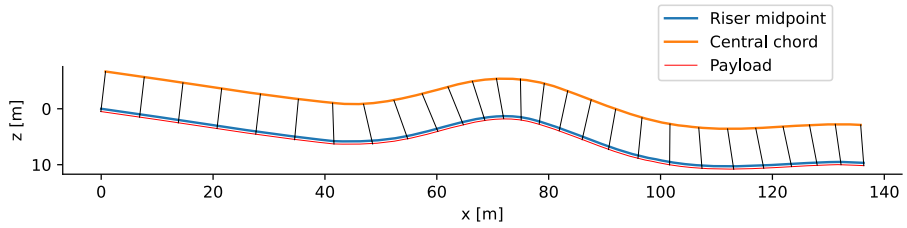

Fig. 9.15: Flight test, rapidly exiting accelerated flight, side view

Black lines are drawn from the riser to the point directly above the payload to help visualize the canopy pitch angle, and are added every 0.5 seconds.

The model predicts the wing configuration will pitch backwards 23° before diving forwards to a pitch angle of -13° which satisfies the expected grading. Although this test is not particularly informative, it's simplicity makes it worthwhile.

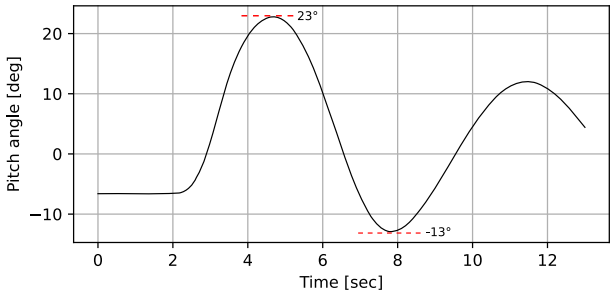

Fig. 9.16: Flight test, rapidly exiting accelerated flight, pitch angle

9.2.3 Steady-state turn

Although the simplicity of longitudinal dynamics make them the best place to start testing a model, the more difficult tests are for the dynamic behavior. One simple test is to check the behavior during a steady 360° maneuver and compare them to the "guidelines" in [13] that lists approximate sink rates and turn radii as a function of bank angle. The method does come with some caveats, however: for example, the author does is not discussing a specific glider, so these values are assumed to be averages of wing performance; this this is a midrange paraglider wing, it is assumed to be "average". Also, the author does not define the control inputs, but standard piloting practice is to use a combination of weight shift and brake for an efficient turn, so it is safe to assume the author is describing situations with those control inputs. Simulating this scenario produces the results in Fig. 9.17:

Table 9.5: Steady-state turn validation

Value	Guideline	Simulated	Error
Turn radius [m]	~12	20	+67%
Sink rate [m/s]	~1.1	1.5	+36%
360° turn rate [sec]	~11.5	16	+40%

Unlike the accurate estimates for the polar curves, which measured steady-state, longitudinal dynamics, this model clearly struggles with this test. It is unclear what is causing

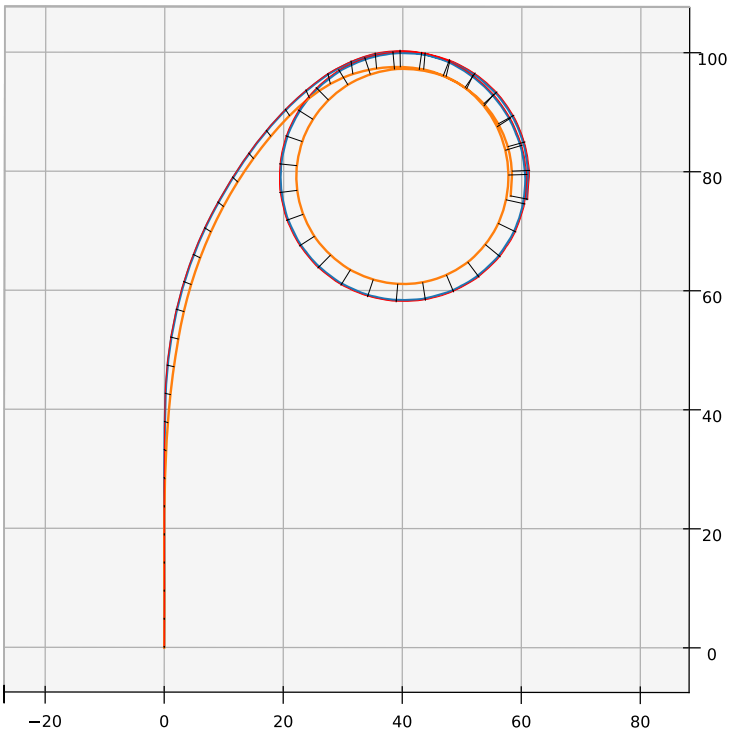

Fig. 9.17: Steady-state turn at a 20° bank angle, top-down view

the discrepancy, but it is an important counterpoint that highlights the many dimensions of model accuracy. It is also suggests a direction for future work on *weight shift modeling*.

Chapter 10

CONCLUSION

10.1 Results

This project completed the set of tasks outline in its *Roadmap*:

1. It developed a novel *Foil geometry* specifically to enable simple representations of paraglider canopies.

2. It selected, implemented and *validated* a fast-but-accurate theoretical aerodynamics model well-suited to the nonlinear geometries and challenging flow conditions of paraglider canopies, as outlined in the *Modeling requirements* defined at the beginning of the project.

3. It developed *parametric models* to estimate the inertial properties and resultant forces of the components of a paraglider.

4. It used the parametric components to *demonstrate* how to produce a complete flight dynamics model of a commercial paraglider wing using only limited technical data, photos, and video of the wing.

5. It *validated* the longitudinal performance of the demonstration model against basic flight test data, as well as highlighted some areas in which the accuracy of flight dynamics could be improved.

This final section of the paper will address the last of the *Modeling requirements*: it will revisit the set of *motivating questions* that helped guide the design process, and consider the ability of these models to answer them.

10.1.1 Study: drag breakdown

A common question for curious pilots is how to reduce the drag of their glider so they can improve the glide ratio or top speed of their wing. The natural progression of this curiosity is wonder where all the drag comes from in the first place. One way to answer that question is to plot the drag contributions from each component [50].

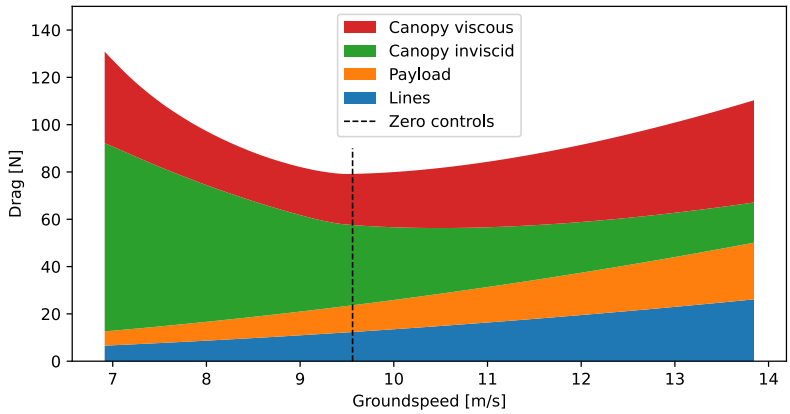

Fig. 10.1: Drag breakdown for Niviuk Hook 3 23 with a pod harness.

Viscous drag includes effects such as the sheer forces produced by the viscosity of the air, and the pressure drag due to flow separation (the "vacuum" that can occur on the downwind side of an object); these forms of drag occur on every surface of the glider, including the lines and payload. Inviscid drag is less intuitive: commonly referred to as "lift-induced drag", it is the energy lost in the vorticity that the wing sheds into its wake as a side-effect of producing lift.

This diagram provides a satisfying look into the behavior of a wing across the range of speeds. At the low end, pilots understand that the "brakes" will slow the wing by increasing its drag, but may be surprised to discover that the increase in drag is dominated by how the wing produces lift. At the high end, it can be surprising to learn what proportion of the total system drag is produced by the seemingly-negligible suspension lines. Although drag is just one piece of the lift/drag ratio, this sort of breakdown is valuable for estimating how much improvement is possible by (for example) reducing the drag of the payload.

This decomposition is also educational because it offers another perspective of how each component of the wing affects the overall design. Consider the general guideline that paraglider wings are designed to achieve their maximum glide ratio at "trim" (zero controls), which usually coincides with the speed that minimizes the total system drag (as seen here). Now suppose the design was changed; for example, increasing the aspect ratio of the canopy will tend to decrease its lift-induced drag, which in turn requires repositioning the payload at trim. The complete system behavior is a complex interaction of components, and having access to a parametric model such as this is an excellent resource for quickly answering questions about glider efficiency by developing an intuition of how their interactions affect the system behavior.

10.1.2 Study: effects of Reynolds numbers and apparent mass

There were two questions at the start of this project that affected my modeling choices:

1. How significant are the effects of apparent mass?

2. How significant are the effects of accurate Reynolds numbers?

Both contributions to the flight dynamics are typically neglected in paraglider dynamics models without clear justification or discussion of their expected impact on model accuracy. The models developed in this paper can be used to provide insight on those questions. Using the Niviuk Hook 3 (size 23) component models created for the *Demonstration*, a programming script created multiple instances of the 6-DoF system models, configuring them to either respect or ignore the effects of apparent mass and precise Reynolds numbers (which are normally computed dynamically for each wing section). Pairs of models — one with the full dynamics and the other lacking one or both effects — are put into a figure-8 maneuver starting at that model's equilibrium state and receiving the same control inputs over a span of 60 seconds. (The maneuver did not use weight shift control to avoid possible issues modeling canopy deformations.) Three simulations were run:

1. To show the affect of neglecting apparent mass (Fig. 10.2)

2. To show the effect of neglecting accurate Reynolds numbers by using a constant $Re = 2 \times 10^6$ (Fig. 10.3)

3. To show the combined effect of neglecting both apparent mass and accurate Reynolds values (Fig. 10.4)

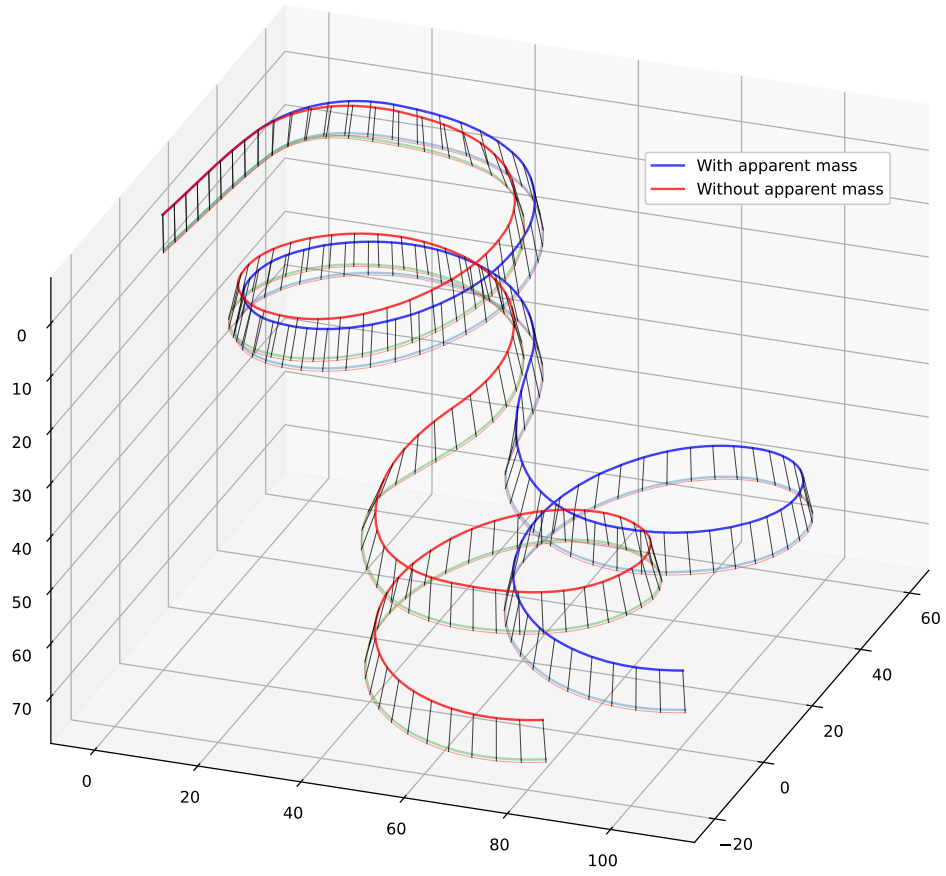

Fig. 10.2: Figure-8 when neglecting apparent mass

The differences produced by each simplification are similar in this case, and will be discussed jointly. First, the less noticeable difference between the two simulations in Fig. 10.4 is the total altitude loss, where the "fixed Reynolds, no apparent mass" model descended an extra 2 meters. The difference is not visually interesting so no side-view is shown, but the effect is worth noting and should be expected for two reasons:

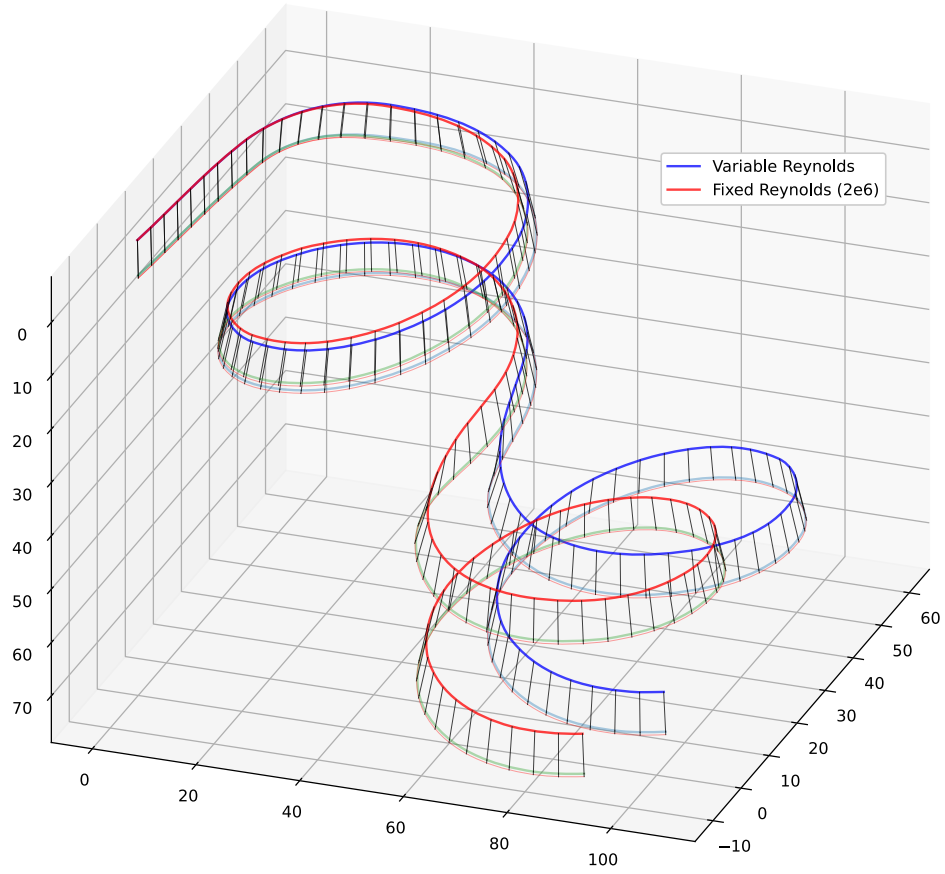

Fig. 10.3: Figure-8 when neglecting accurate Reynolds numbers

1. There is minimal acceleration in the z-direction so the z-component of the apparent mass is negligible.

2. The sections most impacted by the incorrect Reynolds values are at the outside of the span. Since the majority of the lift is produced by the central sections, which are already near the $Re = 2 \times 10^6$ value, total lift is not greatly affected by assuming a fixed value of Re.

The more significant effect was on the lateral motion of the glider, which is easier to see from a top-down perspective (Fig. 10.5), where the complete model exhibited a turn radius of $54\,[m]$ versus $51\,[m]$ of the simplified model. (The cumulative horizontal distances traveled were $522\,[m]$ at $8.7\,\left[\frac{m}{s}\right]$ and $532\,[m]$ at $8.87\,\left[\frac{m}{s}\right]$, respectively.) Again, the effect is expected for two reasons:

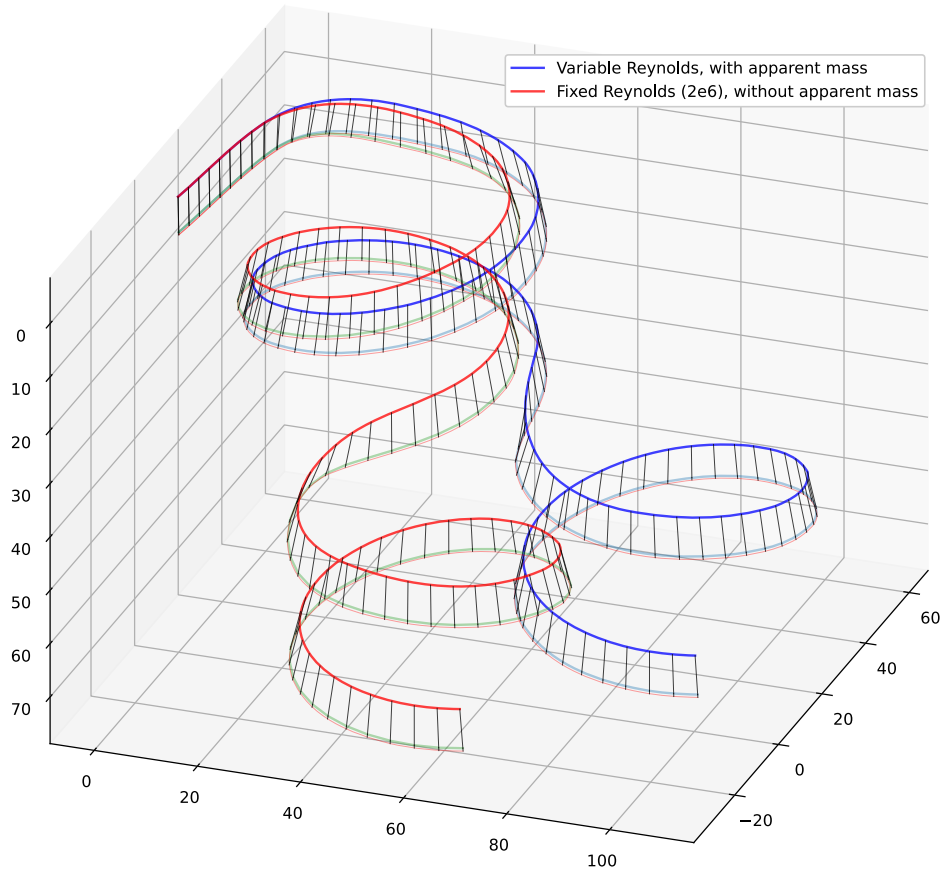

Fig. 10.4: Figure-8 neglecting both apparent mass and accurate Reynolds numbers

1. Apparent mass resists changes to the translational velocity, which reduced the complete models centripetal acceleration and prevented it from producing as narrow a turn as the simplified model.

2. Lower Reynolds values resulted in lower lift coefficients, especially for sections with deflected trailing edges (since their increased curvature magnifies the viscous effects). The lift vectors of sections on the inside semispan are angled into the turn and pull the canopy into the circle, so reducing their lift contributions further reduced the complete models centripetal acceleration.

Because these affects are heavily dependent on the glider design and specific flight maneuvers, this discussion focused on the qualitative nature of these effects. Whether

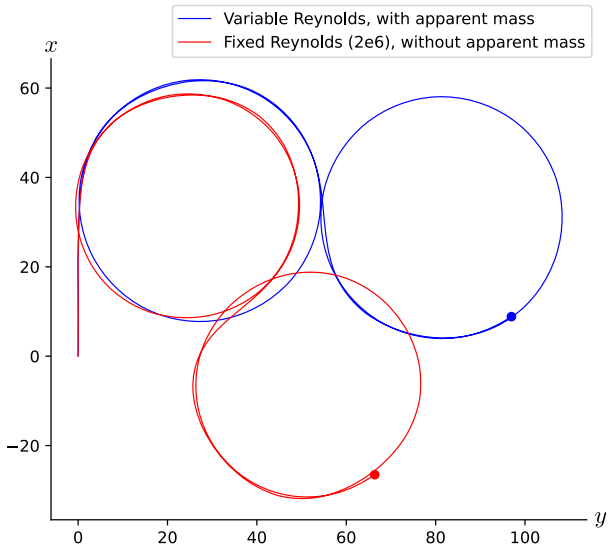

Fig. 10.5: Figure-8 neglecting both apparent mass and accurate Reynolds numbers, topdown view

these sources of error are significant depend heavily on the model (the canopy geometry in particular, as well as target airspeed of the glider) and its application. For example, when developing a linearized model to generate an error term for a control model these effects can be safely neglected, but any long-run simulation should review their specific control sequence (because turning magnifies their impact). With this model, checking the impact of such choices is readily available.

10.1.3 Study: indirect thermal interactions

A reliable way to start a lively discussion on a paragliding forum is to question what happens when a wing encounters a thermal on only one side of its wing. Some pilots will argue that the thermal will pull the wing in; other pilots will argue that the thermal will push the wing away. A grand desire of this project was that the resulting flight dynamics model might be able to shed light on why two seasoned pilots might hold such opposing views.

This final study used the Niviuk Hook 3 size 23 components from the *Demonstration* with a 6-DoF system dynamics model. The scenario is simple: place a thermal slightly off-center of the path of a paraglider flying straight forward at equilibrium with symmetric

brakes. Because the span of the wing is only 8.84 $[m]$, the thermal was placed 15 $[m]$ to the right with exponential falloff such that the thermal strength was reduced to 5% by the time it reached the center of the canopy with a peak (core) strength of 3 $[\frac{m}{s}]$ (extremely strong for such a tight thermal). The effect of the exponential falloff was a peak gradient of 0.67 $[\frac{m}{s}]$ from the wingtip nearest the thermal to the center of the canopy as the glider passed the core.

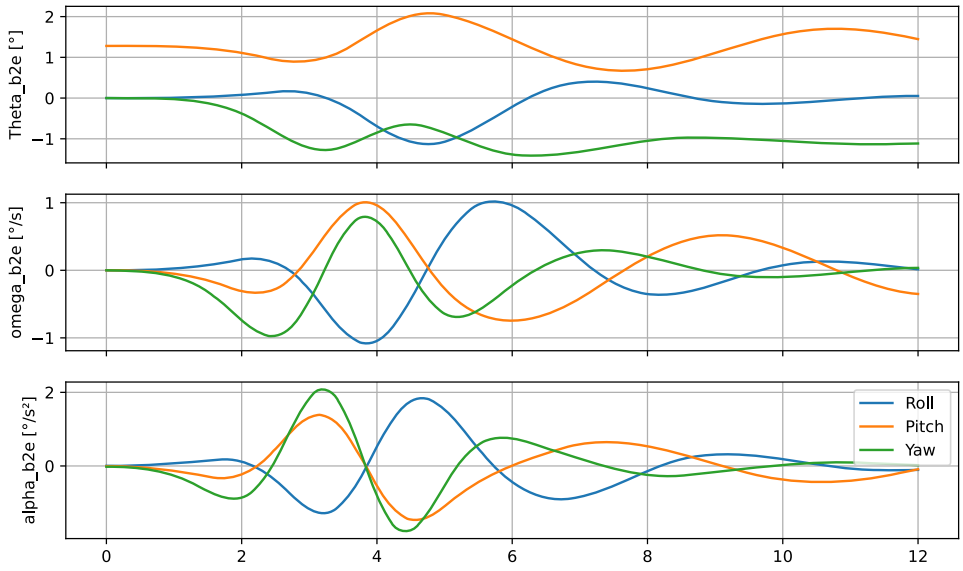

Fig. 10.6: Indirect thermal interaction.

The first row represents the Euler angles for position, the second row represents the angular velocities, and the third row is the angular accelerations.

These results can be viewed in two ways: quantitatively and qualitatively. From a quantitative perspective the results are disappointing: the absolute angular deviations were on the order of 1°, which seem impossibly small for pilots to argue over. From a qualitative perspective, however, the results are perhaps more interesting. As the wing passes the thermal, the canopy initially rolls to the right (into the thermal), pitches forward (into the thermal), and the adverse yaw twists the wing to the left (away from the thermal); although the angular deviations are tiny it may produce an effect similar to falling, which needs only a small distance to produce a striking sensation. The same logic applies after the initial response, where the accelerates again, but more rapidly, and in the opposite direction: now

the wing is rolling away from the thermal while yawing into it. Perhaps the sensation of acceleration holds the key to the argument: whether a pilot is more sensitive to roll or yaw, and whether they're more sensitive to the initial or secondary accelerations may offer a partial explanation?

Personally I find this argument unconvincing. Despite the potential explanation offered by the qualitative analysis, it seems much more likely that the model has failed to capture one or more of the significant dynamics of the system. One possible cause is the foil aerodynamics model, which is not intended to capture unsteady aerodynamics; despite its accuracy in the wind tunnel testing, it may be inadequate for this level of subtlety in dynamic scenarios. Another possible cause is the quasi-rigid-body assumption imposed on the canopy geometry; real wings would flex and distort, especially in such a strong thermal, and it seems like that such deformations may play a larger roll that anticipated.

All in all, despite the underwhelming results the truth is this was always an ambitious goal, and I hope it demonstrates the theoretical advantages of pursuing flight dynamics models that are capable of capturing the effects of non-uniform wind vectors along the span of the wing, and will serve as a starting point for some future work. Perhaps we will someday have an answer for the forums.

10.2 Future work

10.2.1 Canopy

- Arc deformations: the *design curves* that define the foil geometry are not required to be constant functions; they can be functions of control inputs, such as weight shift. The primary difficulty is that the current implementation of the *NLLT* assumes that the shape of the canopy is constant, but that a practical limitation, not a theoretical one.

- Weight shift modeling: the *Steady-state turn* sanity check of the demonstration model suggests that lateral movement of the mass centroid is not the primary control mech-

anism for weight shift control. The alternative mechanism is the wing deformations that occur during weight shift. At the outset of this project the assumption was that the canopy deformations during weight shift would be negligible compared to the displacement of payload mass, but the turn radius and sink rate suggest otherwise. It may be fruitful to generate plausible $yz(s, \delta_w)$ design curves (so the foil arc deforms as a function of weight shift), and consider if the changes to the canopy aerodynamics would explain the inaccuracies in the rigid canopy model. If canopy arc deflections prove to be a significant factor for accurate weight shift predictions, they should probably be implemented as an interaction between $yz(s)$ and the suspension line model. (Paraglider pilots quickly discover the relationship between chest riser strap width and weight shift control, which strongly suggests that the lines play a dominant role).

- Choice of airfoil: the *Demonstration* chose the NACA 24018 as an example of a conservative guess, but if a few commercial section profiles were measured accurately (including their spanwise variation), all models of commercial paraglider wings would benefit.

- Deflected profiles: the demonstration used section *Profiles* produced by a "two circle" model of trailing edge deflection. That optimistic model was designed to balance the accuracy of profile deformation against the ability to estimate the aerodynamic coefficients with XFOIL. In reality, their unnaturally smooth curvature likely causes them to underestimate flow separation. Future work would benefit from more accurate deflection profiles.

- Aerodynamic coefficients: in conjunction with more accurate deflection profiles, another improvement would be is to use more sophisticated methods to estimate the aerodynamic coefficients. One option is RFOIL from Delft University of Technology (a fork of XFOIL that is reported to improve estimates, particularly at high angles of attack), or to apply a complete computational fluid dynamics approach with OpenFoam.

10.2.2 Lines

- The parameters for the *brakes* are confusing at first glance, and tedious to tune. At the least they would benefit from an automated procedure where instead of having to tune $s_{\text{start},1}$ and $s_{\text{stop},1}$ to match κ_b (which was in turn limited by the $\bar{\delta}_{d\max}$ supported by the aerodynamic coefficient set). It would be much easier to define $s_{\text{start},1}$ and $s_{\text{stop},1}$ at some hypothetical value of κ_b and have the lines adjust their values based on the true κ_b.

10.2.3 Harness

- The *spherical model* neglects pitch and yaw moments due to angle of attack and sideslip, but because paragliders put their legs out in front those effects seem likely.

- The harness model uses constant drag coefficients. [20] developed a model for the harness that accounts for Reynolds numbers, but that model was not tested in this work.

10.2.4 System dynamics

- This paper derived a *9-DoF* system dynamics model that modeled the connection between the lines and payload as a spring-damper system, but without flight testing the parameters were difficult to estimate. It would be interesting to review the applicability of the spring-damper model and to estimate suitable parameters. I suspect that the lack of canopy deformations and the inability of the 6-DoF to show payload-relative roll are at least partial explanation of the underwhelming results of the indirect thermal study. The sensation of payload-relative roll and yaw accelerations could definitely play a role in why pilots disagree on the behavior of a paraglider encountering a thermal.

10.3 Open source

The materials to produce this paper and its implementation [1] are both available under permissive open source licenses. Although this work focused on paragliders, the structure of the models is mirrored in the structure of the code, and should be easily adaptable to other gliding aircraft such as hang gliders or kites. For maximum versatility and approachability, the entire implementation was built on the Python scientific computing stack; despite not producing the fastest implementation, Python made up for the performance cost with value in other areas:

- Free (unlike MATLAB, AutoCAD, etc)
- Extensive cross-domain usage (aerospace, computer science, etc)
- Powerful scientific computing libraries (NumPy, SciPy, Numba)
- Easy to integrate into tools with native Python interpreters (such as FreeCAD, Blender, and QGIS)

I am grateful for the work freely shared by those who came before, and hope that this material may provide some value to those who follow.

BIBLIOGRAPHY

[1] Peter F. Heatwole. Glidersim. v2022.03.0a. doi:10.5281/zenodo.6383053.

[2] Sam Davey, Neil Gordon, Ian Holland, Mark Rutten, and Jason Williams. *Bayesian Methods in the Search for MH370*. SpringerBriefs in Electrical and Computer Engineering. Springer Publishing Company, Incorporated, 1st edition, 2016. ISBN 978-981-10-0379-0. doi:10.1007/978-981-10-0379-0.

[3] Snorri Gudmundsson. *General Aviation Aircraft Design*. Butterworth-Heinemann, Boston, January 2014. ISBN 978-0-12-397308-5. doi:10.1016/B978-0-12-397308-5.00001-5.

[4] John David Anderson. *Aircraft Performance and Design*. Volume 1. WCB/McGraw-Hill Boston, 1999.

[5] John Anderson. Ludwig Prandtl's Boundary Layer. *Physics Today*, 58(12):42–48, December 2005. doi:10.1063/1.2169443.

[6] John J. Bertin and Russel M. Cummings. *Aerodynamics for Engineers*. Pearson Education, sixth edition, 2014.

[7] John D. Anderson. *Fundamentals of Aerodynamics*. McGraw-Hill Education, sixth edition, 2017. ISBN 978-1-259-12991-9.

[8] Mark Drela. *Flight Vehicle Aerodynamics*. The MIT Press, February 2014. ISBN 978-0-262-52644-9.

[9] Russel M. Cummings, William H. Mason, Scott A. Morton, and David R. McDaniel. *Applied Computational Aerodynamics*. Cambridge University Press, April 2015. ISBN 978-1-316-23651-2.

[10] Doug McLean. *Understanding Aerodynamics: Arguing from the Real Physics*. John Wiley & Sons, 2012. ISBN 978-1-119-96751-4.

[11] Peter C. Hughes. *Spacecraft Attitude Dynamics*. Dover Publications, Inc., Mineola, New York, 2004. ISBN 978-0-486-43925-9.

[12] Brian L. Stevens, Frank L. Lewis, and Eric N. Johnson. *Aircraft Control and Simulation: Dynamics, Controls Design, and Autonomous Systems*. John Wiley & Sons, 2015. ISBN 978-1-118-87098-3.

[13] Dennis Pagen. *The Art of Paragliding*. Black Mountain Books, 2001. ISBN 0-936310-14-6.

[14] Chris Madsen, Ronald Sostaric, and Chris Cerimele. Flight performance, aerodynamics, and simulation development for the X-38 parafoil test program. In *17th AIAA Aerodynamic Decelerator Systems Technology Conference and Seminar*, 2108. 2003.

[15] Nathan Slegers and Mark Costello. Aspects of control for a parafoil and payload system. *Journal of Guidance, Control, and Dynamics*, 26(6):898–905, 2003. doi:10.2514/2.6933.

[16] John D Nicolaides. Parafoil wind tunnel tests. Technical Report, Notre Dame Univ in Dept of Aerospace and Mechanical Engineering, 1971.

[17] Miguel A. González. Prandtl theory applied to paraglider aerodynamics. In *Aerospace Design Conference*, 1220. 1993. doi:10.2514/6.1993-1220.

[18] Gil Iosilevskii. Lifting-line theory of an arched wing in asymmetric flight. *Journal of Aircraft*, 33(5):1023–1026, 1996. doi:10.2514/3.47050.

[19] Hervé Belloc. Wind tunnel investigation of a rigid paraglider reference wing. *Journal of Aircraft*, 52(2):703–708, 2015. doi:10.2514/1.C032513.

[20] Robert Kulhánek. Identification of a degradation of aerodynamic characteristics of a paraglider due to its flexibility from flight test. *Aircraft Engineering and Aerospace Technology*, January 2019. doi:10.1108/AEAT-06-2018-0162.

[21] Warren F. Phillips and Deryl O. Snyder. Modern adaptation of Prandtl's classic lifting-line theory. *Journal of Aircraft*, 37(4):662–670, 2000. doi:10.2514/2.2649.

[22] Horace Lamb. *Hydrodynamics*. Dover Publications, Inc., New York, USA, June 1945.

[23] Peter Lissaman and Glen J. Brown. Apparent mass effects on parafoil dynamics. In *Aerospace Design Conference*, 1236. 1993. doi:10.2514/6.1993-1236.

[24] Timothy M. Barrows. Apparent mass of parafoils with spanwise camber. *Journal of Aircraft*, 39(3):445–451, 2002. doi:10.2514/2.2949.

[25] Peter G. Thomasson. Equations of motion of a vehicle in a moving fluid. *Journal of Aircraft*, 37(4):630–639, 2000. doi:10.2514/2.2645.

[26] Chiara Toglia and Marilena Vendittelli. Modeling and motion analysis of autonomous paragliders. *Department of computer and system sciences Antonio Ruberti technical reports*, 2010.

[27] Oleg Yakimenko. On the development of a scalable 8-DoF model for a generic parafoil-payload delivery system. In *18th AIAA Aerodynamic Decelerator Systems Technology Conference and Seminar*, 1665. 2005.

[28] Christelle Cumer, Clément Toussaint, Thierry Le Moing, Eric Poquillon, and Yves Coquet. Simulation of generic dynamics flight equations of a parafoil/payload system. In *Control & Automation (MED), 2012 20th Mediterranean Conference On*, 222–228. IEEE, 2012.

[29] J Stephen Lingard. Ram-air parachute Design. In *13th AIAA Aerodynamic Decelerator Systems Technology Conference*, 15–18. 1995.

[30] Holger Babinsky. The aerodynamic performance of paragliders. *The Aeronautical Journal (1968)*, 103(1027):421–428, 1999. doi:10.1017/S0001924000027974.

[31] Diego Muniz Benedetti. *Paragliders Flight Dynamics*. PhD thesis, Universidade Federal de Minas Gerais, 2012.

[32] Pere Casellas. Paraglider Design Handbook. http://laboratoridenvol.com/paragliderdesign/.

[33] Douglas F. Hunsaker and Deryl O. Snyder. A lifting-line approach to estimating propeller/wing interactions. In *24th AIAA Applied Aerodynamics Conference*. San Francisco, CA, USA, 2006. doi:10.2514/6.2006-3466.

[34] Nathan J Slegers. Effects of canopy-payload relative motion on control of autonomous parafoils. *Journal of guidance, control, and dynamics*, 33(1):116–125, 2010.

[35] Chrystine M. Gorman and Nathan Slegers. Evaluation of Multibody Parafoil Dynamics Using Distributed Miniature Wireless Sensors. *Journal of Aircraft*, 49(2):546–555, 2012. doi:10.2514/1.C031566.

[36] Tom Lolies, Nicolas Gourdain, Miguel Charlotte, Herve Belloc, and Bruce Goldsmith. Numerical Methods for Efficient Fluid–Structure Interaction Simulations of Paragliders. *Aerotecnica Missili & Spazio*, 98(3):221–229, September 2019. doi:10.1007/s42496-019-00017-2.

[37] Ira H. Abbott and Albert E. Von Doenhoff. *Theory of Wing Sections*. Dover Publications, Inc., New York, USA, 1959. ISBN 978-0-486-60586-9.

[38] D. Owens. Weissinger's model of the nonlinear lifting-line method for aircraft design. In *36th AIAA Aerospace Sciences Meeting and Exhibit*, Aerospace Sciences Meetings. American Institute of Aeronautics and Astronautics, January 1998. doi:10.2514/6.1998-597.

[39] P. G. Saffman. *Vortex Dynamics*. Cambridge University Press, California Institute of Technology, 1992.

[40] J. Weissinger. The lift distribution of swept-back wings. Technical 1120, NACA, March 1947.

[41] Mark Drela and Harold Youngren. Athena Vortex Lattice.

[42] John D. Anderson and Stephen Corda. Numerical lifting line theory applied to drooped leading-edge wings below and above stall. *Journal of Aircraft*, 17(12):898–904, 1980. doi:10.2514/3.44690.

[43] George M. Ware and James L. Hassell. Wind-tunnel investigation of ram-air inflated all flexible wings of aspect ratios 1.0 to 3.0. Technical Report NASA TM SX-1923, National Aeronautics and Space Administration, 1969.

[44] Nuno Filipe Esperança Virgílio. Study of the aerodynamic efficiency of Free flight equipment. 2004.

[45] Sarah Becker and Paul Bruce. Experimental Study of Paraglider Aerodynamics. Master's thesis, Imperial College London, September 2017. doi:10.13140/RG.2.2.33674.16321.

[46] Maurizio Boffadossi and Federico Savorgnan. Analysis on aerodynamic characteristics of a paraglider airfoil. *AEROTECNICA, MISSILI E SPAZIO*, 95(4):211–218, 2016. doi:10.19249/ams.v95i4.286.

[47] Mark Drela. XFOIL: An Analysis and Design System for Low Reynolds Number Airfoils. In Thomas J. Mueller, editor, *Low Reynolds Number Aerodynamics*, 1–12. Springer Berlin Heidelberg, 1989. doi:10.1007/978-3-642-84010-4_1.

[48] Peter Wild. Airworthiness requirements for hanggliders and paragliders. 2009.

[49] Jewel B. Barlow, William H. Rae, and Alan Pope. *Low-Speed Wind Tunnel Testing*. John Wiley & Sons, third edition, February 1999. ISBN 978-0-471-55774-6.

[50] Holger Babinsky. Aerodynamic improvements of paraglider performance. In *17th Applied Aerodynamics Conference*, 3148. 1999.

[51] Jonathan Blow and Atman J. Binstock. How to find the inertia tensor (or other mass properties) of a 3D solid body represented by a triangle mesh. http://www.number-none.com/blow/inertia/index.html, 2004.

[52] Cha Zhang and Tsuhan Chen. Efficient feature extraction for 2D/3D objects in mesh representation. In *Proceedings 2001 International Conference on Image Processing (Cat. No.01CH37205)*, volume 3, 935–938 vol.3. October 2001.

APPENDIX A

Notation and Symbols

Table A.1: Common Notation

Notation	Meaning
x	a scalar
$\boldsymbol{x}$	a vector
x^y	a scalar raised to a power, where y is a scalar
$\boldsymbol{x}^c$	a vector in the coordinate system c
$\boldsymbol{x}_{B/A}$	a vector from point A to point B ("B with respect to A")
${}^r\dot{\boldsymbol{x}}$	the derivative of a vector taken in reference frame $\mathcal{F}_r$
x_k	a variable at index k of a sequence of length K
$x^{(n)}$	element n of a set of N elements
$\boldsymbol{X}_{M\times N}$	a matrix with M rows and N columns
$\boldsymbol{X}^z$	a matrix exponential, where z is a scalar
$\lvert x\rvert$	absolute value of a scalar
$\lVert\boldsymbol{x}\rVert$	Euclidean norm of a vector
$\lvert\boldsymbol{X}\rvert$	determinant of a matrix
$\boldsymbol{C}_{b/a}$	the directed cosine matrix that transforms vectors from coordinate system a into coordinate system b
$\boldsymbol{q}_{b/a}$	a quaternion that encodes the relative orientation of coordinate system b relative to coordinate system a
$\boldsymbol{\omega}_{b/a}$	angular velocity vector of frame $\mathcal{F}_b$ with respect to frame $\mathcal{F}_a$
$f(\cdot)$, $func(\cdot)$, etc	functions, where `f`, `func`, can be any identifier

Another notation which is useful when building systems of equations involving matrices is the *cross-product matrix operator*, so that $[\boldsymbol{v}]^{\times}\,\boldsymbol{x} \equiv \boldsymbol{v} \times \boldsymbol{x}$:

$$[\boldsymbol{v}]^{\times} \stackrel{\text{def}}{=} \begin{bmatrix} 0 & -v_3 & v_2 \\ v_3 & 0 & -v1 \\ -v_2 & v_1 & 0 \end{bmatrix}$$

By their nature, vectors require the most intricate notation, since a fully specified vector might include all of:

1. A reference frame
2. A coordinate system
3. A fixed point (if it's a bound vector)

For simplicity, Table A.1 only shows examples of each distinct element of a vector encoding. In practice, vectors may appear quite complex; for some realistic examples taken from [12]:

$\boldsymbol{p}_{A/B} \stackrel{\text{def}}{=}$ the position of the point A with respect to point B

$\boldsymbol{v}_{A/i} \stackrel{\text{def}}{=}$ the velocity vector of a point A in frame $\mathcal{F}_i$

${}^{b}\dot{\boldsymbol{v}}_{A/i} \stackrel{\text{def}}{=}$ the vector derivative of $\boldsymbol{v}_{A/i}$ taken in frame $\mathcal{F}_b$

$\boldsymbol{v}^{c}_{A/i} \stackrel{\text{def}}{=}$ array of components of $\boldsymbol{v}_{A/i}$ in coordinate system c

${}^{b}\dot{\boldsymbol{v}}^{c}_{A/i} \stackrel{\text{def}}{=}$ components in coordinate system c of the derivative taken in frame $\mathcal{F}_b$

APPENDIX B

Derivations

B.1 Parametric design curves

The *"simplified" foil geometry* chose a set of variables (3.15) that describe different aspects of the shape. This section provides definitions for several of those variables using parametric functions that can approximate the structure of a typical parafoil using a small number of simple parameters.

B.1.1 Elliptical chord

A *Foil geometry* requires a chord distribution $c(s)$. For parafoils, the chords lengths are most commonly defined by a truncated elliptical function of section index, in which case the distribution is a function of two design parameters. The typical choices are either the root and wingtip chord lengths, or the root length and a taper ratio. Choosing the root and wingtip chord lengths, a truncated elliptical function over the section index $-1 \leq s \leq 1$ is then:

$$\begin{aligned} a &= \frac{1}{\sqrt{1 - \left(\frac{c_{\text{tip}}}{c_{\text{root}}}\right)^2}} \\ b &= c_{\text{root}} \\ c(s) &= b\sqrt{1 - \left(\frac{s}{a}\right)^2} \end{aligned} \tag{2.1}$$

Refer to `EllipticalChord` in `glidersim` for an implementation.

B.1.2 Elliptical arc

In this paper the *arc* of a parafoil is the vector-valued function of $\langle y, z \rangle$ coordinates that position the section reference points. For parafoils, the arc is typically defined by an elliptical function.

A centered elliptical curve can be defined as a function of four parameters, but the symmetry of the wing reduces that to three free design parameters, and normalizing the

arc length reduces it to just two. There are several possible parametrizations, but an intuitive choice is the mean anhedral angle Γ_{tip} and the section roll angle ϕ_{tip} of the wing tips [31].

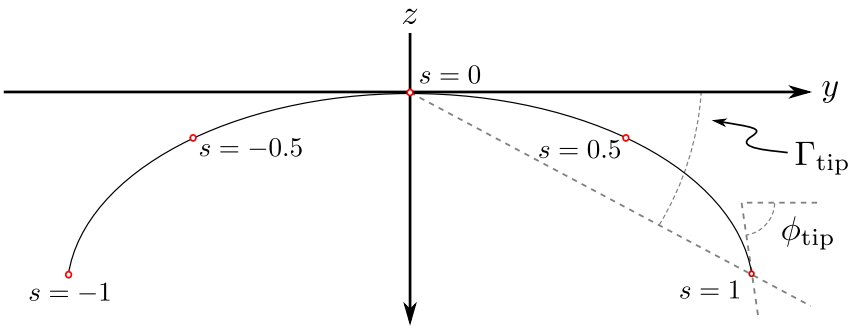

Fig. B.1: Parametrized elliptical arc

Choosing those parameters to define an elliptical function that is proportional to the desired yz-curve produces an intermediate result:

$$k_1 = 1 - \frac{\tan(\Gamma_{\text{tip}})}{\tan(\phi_{\text{tip}})}$$
$$k_2 = 1 - \frac{2\tan(\Gamma_{\text{tip}})}{\tan(\phi_{\text{tip}})}$$
$$A = \frac{k_1}{\sqrt{k_2}}$$
$$B = \frac{k_1}{k_2}\tan(\Gamma_{\text{tip}})$$
$$\boldsymbol{f}(t) = \langle A\cos(t), B\sin(t)\rangle$$

This design requires that $\phi_{\text{tip}} > 2\Gamma_{\text{tip}}$ (so the wing must be wider than it is tall and the wing tip roll cannot exceed 90°) and is valid over $t_{min} \le t \le \pi - t_{min}$, where $t_{min} = \arccos\left(\frac{1}{A}\right)$.

Next although the shape produced by this intermediate result is proportional to the desired curve, it is not directly usable by the *Foil geometry*. It needs two modifications:

1. Make the arc a function of the chosen section index s
2. Scale the arc to a total curve length of b_{flat}

Both can be achieved by normalizing the elliptical function to a curve length of 2. First, scale the axes to produce a new semi-ellipse with a total curve length of 1:

$$L(t) = \int_{\frac{\pi}{2}}^{t_{min}} \|\boldsymbol{f}(t)\| \, dt$$

$$k_3 = L(t_{min})$$

$$\bar{\boldsymbol{f}}(t) = \left\langle \frac{A}{k_3}\cos(t), \frac{B}{k_3}\sin(t) \right\rangle$$

The fact that the simplified foil geometry chose to define the *section index* s as the linear distance along the yz-curve enables a convenient conversion over $\frac{\pi}{2} \leq t \leq t_{min}$ and $0 \leq s \leq 1$:

$$\bar{L}(t) = \int_{\frac{\pi}{2}}^{t_{min}} \left\|\bar{\boldsymbol{f}}(t)\right\| \, dt = s(t)$$

$$t(s) = s^{-1}(t)$$

Thus the complete parametric function for the yz-curve of the arc is thus $\langle y, z \rangle (s) = \bar{\boldsymbol{f}}(t(|s|))$. The integrals and inverse functions are not available analytically, but are trivial to compute numerically. Refer to `EllipticalArc` in `glidersim` for an implementation.

B.1.3 Polynomial torsion

Like most wings, parafoils use section-relative pitch $\theta(s)$ (conventionally referred to as *geometric torsion*) to fine-tune wing behavior. The exact distribution of geometric torsion along a wing can be difficult to measure, but they are frequently described using simple polynomials or piecewise-linear functions. For idealized models of nonlinear geometries such as those developed here, a piecewise-polynomial function is assumed to be adequate.

Assuming a symmetric wing, define three parameters:

- T: the maximum torsion (in radians) at the wingtips
- s_{start}: the section index where the torsion begins (where $0 \leq s_{start} < 1$)
- β: the degree of the polynomial (for example, $\beta = 1$ is linear, $\beta = 2$ is quadratic, etc.)

$$p(s) = \frac{|s| - s_{start}}{1 - s_{start}}$$

$$\theta(s) = \begin{cases} 0 & |s| < s_{start} \\ Tp^{\beta} & |s| \geq s_{start} \end{cases}$$

Refer to `PolynomialTorsion` in `glidersim` for an implementation.

B.2 Area and Volume of a Mesh

The paraglider dynamics require the inertial properties of the canopy surface areas and volume. These include the magnitudes (total mass or volume), centroids, and inertia tensors. All of these quantities can be computed using a triangular surface mesh over the canopy surfaces.

What follows is a reproduction of the procedure developed in [51], which is a functionally equivalent to the procedure from [52] but with a more intuitive interpretation and complete equations for the inertia tensors.

B.2.1 Area

First, for each of the upper and lower surfaces, cover the surface with a triangulated mesh so it is represented by a set of N triangles. Each triangle is defined by three points $\{\mathrm{P1}, \mathrm{P2}, \mathrm{P3}\}_n$ in canopy coordinates. For convenience, define position vectors for each of the three points of the nth triangle: $\boldsymbol{r}_{i,n} \stackrel{\text{def}}{=} \boldsymbol{r}_{Pi/O,n}$.

The area of each triangle is easily computed using the vector cross-product of two legs of the triangle:

$$a_n = \frac{1}{2}\rho \left\| (\boldsymbol{r}_{2,n} - \boldsymbol{r}_{1,n}) \times (\boldsymbol{r}_{3,n} - \boldsymbol{r}_{2,n}) \right\|$$

The total area of the surface is the sum of the triangle areas:

$$a = \sum_{n=1}^{N} a_n$$

The area centroid of each triangle:

$$\overline{\boldsymbol{a}}_n \stackrel{\text{def}}{=} \frac{1}{3}\left(\boldsymbol{r}_{1,n} + \boldsymbol{r}_{2,n} + \boldsymbol{r}_{3,n}\right)$$

And the centroid A of the total surface area with respect to the canopy origin O:

$$\boldsymbol{r}_{\mathrm{A/O}} = \frac{1}{a}\sum_{n=1}^{N} a_n \overline{\boldsymbol{a}}_n$$

The covariance matrix of the total surface area:

$$\boldsymbol{\Sigma}_a = \sum_{n=1}^{N} a_n \overline{\boldsymbol{a}}_n \overline{\boldsymbol{a}}_n^T$$

The inertia tensor of the total surface area a about the canopy origin O:

$$\boldsymbol{J}_{a/\mathrm{O}} = \mathrm{trace}\left(\boldsymbol{\Sigma}_a\right)\boldsymbol{I}_3 - \boldsymbol{\Sigma}_a$$

This completes the calculation of the three relevant properties for each surface area: the total area a, the area centroid $\boldsymbol{r}_{\mathrm{A/O}}$, and the inertia tensor $\boldsymbol{J}_{a/\mathrm{O}}$.

B.2.2 Volume

Now for the volume. For the purposes of computing the inertia properties of the enclosed air, it is convenient to neglect the air intakes and treat the canopy as a closed volume. Given this simplifying assumption, build another triangular mesh that covers the entire canopy surface as well as the left and right wing tip sections. For this derivation, it is essential that the points on each triangle are ordered such that a right-handed traversal produces a normal vector pointing out of the volume. It is also essential that the complete mesh does not contain any holes, or the volume may be miscounted. Given a surface triangulation over the closed canopy geometry using N triangles, the volume can be computed as follows.

First, treat each triangle as the face of a tetrahedron that includes the origin. The signed volume of the tetrahedron formed by each triangle is given by:

$$v_n = \frac{1}{6}\left(\boldsymbol{r}_{1,n} \times \boldsymbol{r}_{2,n}\right) \cdot \boldsymbol{r}_{3,n}$$

Given that the vertices of each triangle were oriented such that they satisfy a right-hand rule, the sign of each volume will be positive if the normal vector for each triangular face points away from the origin, and negative if it points towards the origin. In essence the tetrahedrons "overcount" the volume for triangles pointing away from the origin, then the

triangles facing the origin subtract away the excess volume. The final volume of the canopy is the simple sum:

$$v = \sum_{n=1}^{N} v_n$$

For the volume centroid of each tetrahedron:

$$\overline{\boldsymbol{v}}_n \stackrel{\text{def}}{=} \frac{1}{4} \sum_{i=1}^{3} \boldsymbol{r}_{i,n}$$

And the centroid V of the total volume with respect to the canopy origin O:

$$\boldsymbol{r}_{\text{V/O}} = \frac{1}{v} \sum_{n=1}^{N} v_n \overline{\boldsymbol{v}}_n$$

Lastly, calculating the inertia tensor of the volume can be simplified by computing the inertia tensor of a prototypical or "canonical" tetrahedron and applying an affine transformation to produce the inertia tensor of each individual volume.

First, given the covariance matrix of the "canonical" tetrahedron:

$$\hat{\boldsymbol{\Sigma}} \stackrel{\text{def}}{=} \begin{bmatrix} \frac{1}{60} & \frac{1}{120} & \frac{1}{120} \\ \frac{1}{120} & \frac{1}{60} & \frac{1}{120} \\ \frac{1}{120} & \frac{1}{120} & \frac{1}{60} \end{bmatrix}$$

Use the points in each triangle to define:

$$\boldsymbol{T}_n \stackrel{\text{def}}{=} \begin{bmatrix} | & | & | \\ \boldsymbol{r}_{1,n} & \boldsymbol{r}_{2,n} & \boldsymbol{r}_{3,n} \\ | & | & | \end{bmatrix}$$

The covariance of each tetrahedron volume is then:

$$\boldsymbol{\Sigma}_n = |\boldsymbol{T}_n| \, \boldsymbol{T}_n^T \hat{\boldsymbol{\Sigma}} \boldsymbol{T}_n$$

And the covariance matrix of the complete volume:

$$\boldsymbol{\Sigma}_v = \sum_{n=1}^{N} \boldsymbol{\Sigma}_n$$

And at last, the inertia tensor of the volume about the origin O can be computed directly from the covariance matrix:

$$\boldsymbol{J}_{v/O} = \text{trace}\left(\boldsymbol{\Sigma}_v\right) \boldsymbol{I}_3 - \boldsymbol{\Sigma}_v$$

B.3 Apparent mass of a parafoil

This section presents Barrows' method [24] for estimating the apparent mass matrix of a wing with circular arc anhedral. (For a discussion of apparent mass effects, see *Apparent Mass*.) The equations have been adapted to use the standard notation of this paper.

The purpose of the equations is estimate several terms that allow the paraglider system dynamics model to calculate the apparent inertia matrix with respect to the dynamics reference point, so the apparent mass can be taken into account when calculating the canopy acceleration. The necessary terms are:

- $\boldsymbol{A}_{a/R}$: apparent inertia matrix with respect to some *reference point* R. This matrix is comprised of a translational inertia part $\boldsymbol{M}_a$ and a rotational inertia part $\boldsymbol{J}_{a/R}$.
- $\boldsymbol{r}_{RC/R}$: roll center with respect to R
- $\boldsymbol{r}_{PC/RC}$: pitch center with respect to the *roll center* RC

Some notes about Barrows' development:

- It assumes the foil is symmetric about the xz-plane (left-right symmetry) and about the yz-plane (fore-aft symmetry).
- It requires that the dynamics reference point R lies in the xz-plane
- It assumes the canopy arc is circular.
- It assumes a constant chord length over the entire span.
- It assumes constant thickness over the entire span.
- It assumes no chordwise camber.
- It assumes the chords are all parallel to the x-axis (which also means no geometric twist). This mostly isn't a problem since our coordinate system is defined by the central chord, the geometric torsion angles tend to be quite small, and twist tends to

occur over segments which represent negligible volume compared to the bulk of the wing.

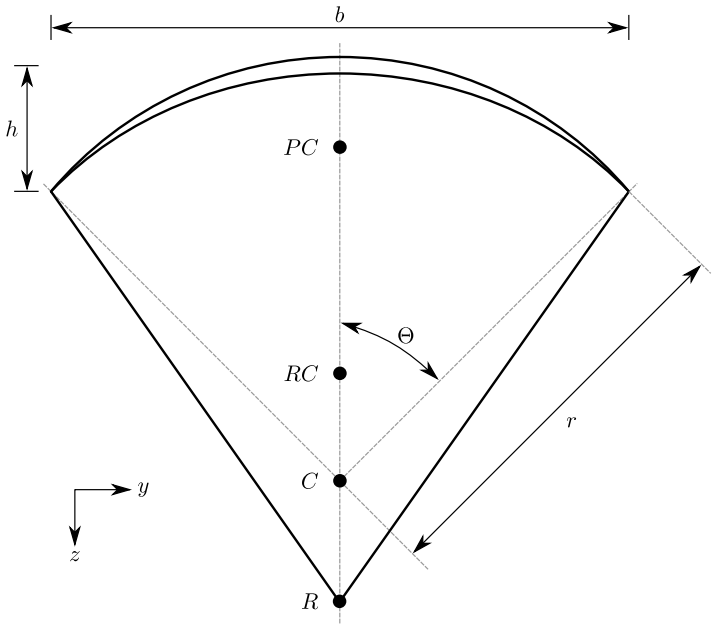

Fig. B.2: Geometry for Barrow's apparent mass equations.

Some initial definitions:

$$t = \text{Airfoil thickness.}$$
$$h^* = \frac{h}{b}$$

First, the apparent mass terms for a flat wing of a similar volume, from Barrows' equations 34-39:

$$m_{f11} = k_A \pi \left(t^2 b/4\right)$$
$$m_{f22} = k_B \pi \left(t^2 c/4\right)$$
$$m_{f33} = \left[\mathrm{AR}/\left(1+\mathrm{AR}\right)\right] \pi \left(c^2 b/4\right)$$

$$I_{f11} = 0.055 \left[\mathrm{AR}/\left(1+\mathrm{AR}\right)\right] b S^2$$
$$I_{f22} = 0.0308 \left[\mathrm{AR}/\left(1+\mathrm{AR}\right)\right] c^3 S$$
$$I_{f33} = 0.055 b^3 t^2$$

Where k_A and k_B are the "correction factors for three-dimensional effects":

$$k_A = 0.85$$

$$k_B = 1.0$$

Assuming the parafoil arc is circular and with no chordwise camber, use Barrows equations 44 and 50 to compute the *pitch center PC* and *roll center RC* as points directly above the *confluence point C* of the arc:

$$z_{PC/C} = -\frac{r \sin(\Theta)}{\Theta}$$

$$z_{RC/C} = -\frac{z_{PC/C}\, m_{f22}}{m_{f22} + I_{f11}/r^2}$$

$$z_{PC/RC} = z_{PC/C} - z_{RC/C}$$

Modifying the apparent mass terms from the flat wing to approximate the terms for the arched wing, Barrows equations 51-55:

$$m_{11} = k_A \left[1 + \left(\frac{8}{3}\right) h^{*2}\right] \pi \left(t^2 b/4\right)$$

$$m_{22} = \frac{r^2 m_{f22} + I_{f11}}{z_{PC/C}^2}$$

$$m_{33} = m_{f33}$$

$$I_{11} = \frac{z_{PC/RC}^2}{z_{PC/C}^2} r^2 m_{f22} + \frac{z_{RC/C}^2}{z_{PC/C}^2} I_{f11}$$

$$I_{22} = I_{f22}$$

$$I_{33} = 0.055 \left(1 + 8h^{*2}\right) b^3 t^2$$

The apparent mass and apparent moment of inertia matrices are then defined in Barrows equations 1 and 17:

$$\boldsymbol{M}_a \stackrel{\text{def}}{=} \begin{bmatrix} m_{11} & 0 & 0 \\ 0 & m_{22} & 0 \\ 0 & 0 & m_{33} \end{bmatrix} \tag{2.2}$$

$$\boldsymbol{I}_a \stackrel{\text{def}}{=} \begin{bmatrix} I_{11} & 0 & 0 \\ 0 & I_{22} & 0 \\ 0 & 0 & I_{33} \end{bmatrix} \tag{2.3}$$

Define two helper matrices:

$$\boldsymbol{S}_2 \stackrel{\text{def}}{=} \begin{bmatrix} 0 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \end{bmatrix}$$

$$\boldsymbol{Q} = \boldsymbol{S}_2 \left[\boldsymbol{r}_{PC/RC}\right]^{\times} \boldsymbol{M}_a \left[\boldsymbol{r}_{RC/R}\right]^{\times}$$

Where $[\boldsymbol{x}]^{\times}$ is the *cross-product matrix operator*.

Using the helper matrices, use Barrows equation 25 to write the rotational part of the apparent inertia matrix:

$$\boldsymbol{J}_{a/R} \stackrel{\text{def}}{=} \boldsymbol{I} - \left[\boldsymbol{r}_{RC/R}\right]^{\times} \boldsymbol{M}_a \left[\boldsymbol{r}_{RC/R}\right]^{\times} - \left[\boldsymbol{r}_{PC/RC}\right]^{\times} \boldsymbol{M}_a \left[\boldsymbol{r}_{PC/RC}\right]^{\times} \boldsymbol{S}_2 - \boldsymbol{Q} - \boldsymbol{Q}^T$$

And the corresponding angular momentum of the apparent mass about R, using Barrows equation 24:

$$\boldsymbol{h}_{a/R} = \left(\boldsymbol{S}_2 \left[\boldsymbol{r}_{PC/RC}\right]^{\times} + \left[\boldsymbol{r}_{RC/R}\right]^{\times}\right) \boldsymbol{M}_a \boldsymbol{v}_{R/e} + \boldsymbol{J}_{a/R}\omega$$

And finally, the completed apparent inertia matrix with respect to the reference point R, from Barrows equation 27:

$$\boldsymbol{A}_{a/R} = \begin{bmatrix} \boldsymbol{M}_a & -\boldsymbol{M}_a \left(\left[\boldsymbol{r}_{RC/R}\right]^{\times} + \left[\boldsymbol{r}_{PC/RC}\right]^{\times} \boldsymbol{S}_2\right) \\ \left(\boldsymbol{S}_2 \left[\boldsymbol{r}_{PC/RC}\right]^{\times} + \left[\boldsymbol{r}_{RC/R}\right]^{\times}\right) \boldsymbol{M}_a & \boldsymbol{J}_{a/R} \end{bmatrix} \tag{2.4}$$

Plus the vectors necessary to incorporate $\boldsymbol{J}_{a/R}$ into the final dynamics:

$$\boldsymbol{r}_{PC/RC} = \begin{bmatrix} 0 & 0 & z_{PC/RC} \end{bmatrix}$$

Linear momentum of the apparent mass:

$$\boldsymbol{p}_{a/e} = \boldsymbol{M}_a \cdot \left(\boldsymbol{v}_{R/e} - \left[\boldsymbol{r}_{RC/R}\right]^{\times} \omega_{b/e} - \left[\boldsymbol{r}_{PC/RC}\right]^{\times} \boldsymbol{S}_2 \cdot \omega_{b/e}\right) \tag{2.5}$$

Angular momentum of the apparent mass about R:

$$\boldsymbol{h}_{a/R} = \left(\boldsymbol{S}_2 \cdot \left[\boldsymbol{r}_{PC/RC}\right]^{\times} + \left[\boldsymbol{r}_{RC/R}\right]^{\times}\right) \cdot \boldsymbol{M}_a \cdot \boldsymbol{v}_{R/e} + \boldsymbol{J}_{a/R} \cdot \omega_{b/e} \tag{2.6}$$

Refer to `ParagliderWing` in `glidersim` for an implementation.

B.4 Paraglider system models

B.4.1 Model 6a

This section describe a paraglider dynamics model with 6 degrees of freedom. It uses a rigid-body assumption, and incorporates the effects of apparent mass. The dynamics are computed with respect to the riser midpoint RM instead of the wing center of mass B because it avoids needing to recompute the apparent inertia matrix whenever B changes. In this derivation all vectors are in the canopy coordinate system c, so the vector coordinate systems are implicit in the notation.

The derivation develops the equations of motion by starting with derivatives of linear and angular momentum. The derivation is largely based on the excellent [11], although this section uses this paper's version of Stevens' notation (see *Notation and Symbols*).

An implementation of this model is available as `Paraglider6a` in the `glidersim` package. The `glidersim` package also includes `Paraglider6b` and `Paraglider6c`, which decouple the translational and angular equations of motion by choosing the glider center of gravity for the dynamics reference point, but do not incorporate the apparent mass matrix.

B.4.1.1 Real mass only Start with the equations for the translational and angular momentum of the body b about the reference point RM as observed by the inertial reference frame e:

$$\begin{aligned} \boldsymbol{p}_{b/e} &= m_b \, \boldsymbol{v}_{B/e} \\ &= m_b \left(\boldsymbol{v}_{RM/e} + \boldsymbol{\omega}_{b/e} \times \boldsymbol{r}_{B/RM}\right) \end{aligned} \tag{2.7}$$

$$\boldsymbol{h}_{b/RM} = m_b \, \boldsymbol{r}_{B/RM} \times \boldsymbol{v}_{RM/e} + \boldsymbol{J}_{b/RM} \cdot \boldsymbol{\omega}_{b/e} \tag{2.8}$$

Compute the momentum derivatives in the inertial frame $\mathcal{F}_e$ in terms of derivatives in the body frame $\mathcal{F}_b$:

$$\begin{aligned}
{}^{e}\dot{\boldsymbol{p}}_{b/e} &= {}^{b}\dot{\boldsymbol{p}}_{b/e} + \boldsymbol{\omega}_{b/e} \times \boldsymbol{p}_{b/e} \\
&= m_b \left({}^{b}\dot{\boldsymbol{v}}_{RM/e} + {}^{b}\dot{\boldsymbol{\omega}}_{b/e} \times \boldsymbol{r}_{B/RM} + \boldsymbol{\omega}_{b/e} \times \cancelto{0}{{}^{b}\dot{\boldsymbol{r}}_{B/RM}} \right) + \boldsymbol{\omega}_{b/e} \times \boldsymbol{p}_{b/e} \\
&= m_b \left({}^{b}\dot{\boldsymbol{v}}_{RM/e} + {}^{b}\dot{\boldsymbol{\omega}}_{b/e} \times \boldsymbol{r}_{B/RM} \right) + \boldsymbol{\omega}_{b/e} \times \boldsymbol{p}_{b/e}
\end{aligned}$$

$$\begin{aligned}
{}^{e}\dot{\boldsymbol{h}}_{b/RM} &= {}^{b}\dot{\boldsymbol{h}}_{b/RM} + \boldsymbol{\omega}_{b/e} \times \boldsymbol{h}_{b/RM} \\
&= m_b \left(\cancelto{0}{{}^{b}\dot{\boldsymbol{r}}_{B/RM}} \times \boldsymbol{v}_{RM/e} + \boldsymbol{r}_{B/RM} \times {}^{b}\dot{\boldsymbol{v}}_{RM/e} \right) + \boldsymbol{J}_{b/RM} \cdot {}^{b}\dot{\boldsymbol{\omega}}_{b/e} + \boldsymbol{\omega}_{b/e} \times \boldsymbol{h}_{b/RM} \\
&= m_b \, \boldsymbol{r}_{B/RM} \times {}^{b}\dot{\boldsymbol{v}}_{RM/e} + \boldsymbol{J}_{b/RM} \cdot {}^{b}\dot{\boldsymbol{\omega}}_{b/e} + \boldsymbol{\omega}_{b/e} \times \boldsymbol{h}_{b/RM}
\end{aligned} \tag{2.9}$$

Relate the derivatives of momentum with respect to the inertial frame to the net force on the body $\boldsymbol{f}_b$ and the net moment on the body about the reference point $\boldsymbol{g}_{b/RM}$:

$$\begin{aligned}
{}^{e}\dot{\boldsymbol{p}}_{b/e} &= \boldsymbol{f}_b \\
{}^{e}\dot{\boldsymbol{h}}_{b/RM} + \boldsymbol{v}_{RM/e} \times \boldsymbol{p}_{b/e} &= \boldsymbol{g}_{b/RM}
\end{aligned} \tag{2.10}$$

Where

$$\boldsymbol{f}_b = \boldsymbol{f}_{b,\text{aero}} + \boldsymbol{f}_{b,\text{weight}}$$

$$\boldsymbol{g}_{b/RM} = \boldsymbol{g}_{b,\text{aero}} + \boldsymbol{r}_{B/RM} \times \boldsymbol{f}_{b,\text{weight}}$$

Combining (2.9) and (2.10) gives the final equations for the dynamics of the real mass (solid mass plus the enclosed air) in terms of ${}^{b}\dot{\boldsymbol{v}}_{RM/e}$ and ${}^{b}\dot{\boldsymbol{\omega}}_{b/e}$.

$$\begin{aligned}
m_b \, {}^{b}\dot{\boldsymbol{v}}_{RM/e} + m_b \, {}^{b}\dot{\boldsymbol{\omega}}_{b/e} \times \boldsymbol{r}_{B/RM} &= \boldsymbol{f}_b - \boldsymbol{\omega}_{b/e} \times \boldsymbol{p}_{b/e} \\
m_b \, \boldsymbol{r}_{B/RM} \times {}^{b}\dot{\boldsymbol{v}}_{RM/e} + \boldsymbol{J}_{b/RM} \cdot {}^{b}\dot{\boldsymbol{\omega}}_{b/e} &= \boldsymbol{g}_{b/RM} - \boldsymbol{\omega}_{b/e} \times \boldsymbol{h}_{b/RM} - \boldsymbol{v}_{RM/e} \times \boldsymbol{p}_{b/e}
\end{aligned} \tag{2.11}$$

Rewriting the equations as a linear system:

$$\boldsymbol{A}_{r/RM} \begin{bmatrix} {}^{b}\dot{\boldsymbol{v}}_{RM/e} \\ {}^{b}\dot{\boldsymbol{\omega}}_{b/e} \end{bmatrix} = \begin{bmatrix} \boldsymbol{b}_1 \\ \boldsymbol{b}_2 \end{bmatrix} \tag{2.12}$$

Where:

$$\boldsymbol{A}_{r/RM} = \begin{bmatrix} m_b \boldsymbol{I}_3 & -m_b \left[\boldsymbol{r}_{B/RM}\right]^\times \\ m_b \left[\boldsymbol{r}_{B/RM}\right]^\times & \boldsymbol{J}_{b/RM} \end{bmatrix}$$

$$\boldsymbol{b}_1 = \boldsymbol{f}_b - \boldsymbol{\omega}_{b/e} \times \boldsymbol{p}_{b/e}$$

$$\boldsymbol{b}_2 = \boldsymbol{g}_{b/RM} - \boldsymbol{\omega}_{b/e} \times \boldsymbol{h}_{b/RM} - \boldsymbol{v}_{RM/e} \times \boldsymbol{p}_{b/e}$$

B.4.1.2 Real mass + apparent mass Writing the dynamics in matrix form not only makes it straightforward to solve for the state derivatives, it also makes it easy to incorporate the apparent inertia matrix from *Apparent mass of a parafoil*. Adding the apparent inertia into the system matrix and accounting for the translational and angular apparent momentum produces:

$$\left[\boldsymbol{A}_{r/RM} + \boldsymbol{A}_{a/RM}\right] \begin{bmatrix} {}^b\dot{\boldsymbol{v}}_{RM/e} \\ {}^b\dot{\boldsymbol{\omega}}_{b/e} \end{bmatrix} = \begin{bmatrix} \boldsymbol{b}_3 \\ \boldsymbol{b}_4 \end{bmatrix} \tag{2.13}$$

$$\boldsymbol{b}_3 = \boldsymbol{b}_1 - \boldsymbol{\omega}_{b/e} \times \boldsymbol{p}_{a/e}$$

$$\boldsymbol{b}_4 = \boldsymbol{b}_2 - \boldsymbol{v}_{RM/e} \times \boldsymbol{p}_{a/e} - \boldsymbol{\omega}_{b/e} \times \boldsymbol{h}_{a/RM} + \boldsymbol{v}_{RM/e} \times \left(\boldsymbol{M}_a \cdot \boldsymbol{v}_{RM/e}\right)$$

Where $\boldsymbol{A}_{a/RM}$ is the apparent inertia matrix of the canopy from (2.4), $\boldsymbol{M}_a$ is the apparent mass matrix from (2.2), and $\boldsymbol{p}_{a/e}$ and $\boldsymbol{h}_{a/RM}$ are the linear and angular apparent momentums from (2.5) and (2.6). The extra term $\boldsymbol{v}_{RM/e} \times \left(\boldsymbol{M}_a \boldsymbol{v}_{RM/e}\right)$ in $\boldsymbol{b}_4$ is necessary to avoid double counting the aerodynamic moment already accounted for by the section pitching coefficients.

B.4.2 Model 6b

Following the same logic as *Model 6a*, but targeting ${}^b\boldsymbol{v}_{B/e}$ and using the momentum about the body center of mass B produces a simpler model with a diagonal system matrix, but at the cost of requiring the body center of mass to be determined before computing the apparent inertia matrix with respect to that point. For that reason the apparent mass is

neglected here, although if B lies in the xz-plane then the method described in *Apparent mass of a parafoil* could be used.

The main purpose of this model is for validating model implementations. An implementation of this model is available as `Paraglider6b` in the `glidersim` package.

$$\boldsymbol{p}_{b/e} = m_b \, \boldsymbol{v}_{B/e} \tag{2.14}$$

$$\boldsymbol{h}_{b/B} = \boldsymbol{J}_{b/B} \cdot \boldsymbol{\omega}_{b/e} \tag{2.15}$$

Computing the inertial derivatives with respect to the body frame:

$$\begin{aligned} {}^{e}\dot{\boldsymbol{p}}_{b/e} &= m_b \, {}^{b}\dot{\boldsymbol{v}}_{B/e} + \boldsymbol{\omega}_{b/e} \times \boldsymbol{p}_{b/e} \\ {}^{e}\dot{\boldsymbol{h}}_{b/B} &= \boldsymbol{J}_{b/B} \cdot {}^{b}\dot{\boldsymbol{\omega}}_{b/e} + \boldsymbol{\omega}_{b/e} \times \boldsymbol{h}_{b/B} \end{aligned} \tag{2.16}$$

Using the body center of mass as the reference point simplifies the equation for angular momentum:

$$\begin{aligned} {}^{e}\dot{\boldsymbol{p}}_{b/e} &= \boldsymbol{f}_b \\ {}^{e}\dot{\boldsymbol{h}}_{b/B} &= \boldsymbol{g}_{b/B} \end{aligned} \tag{2.17}$$

Combining (2.16) and (2.17): and rewriting as a linear system:

$$\begin{bmatrix} m_b & 0 \\ 0 & \boldsymbol{J}_{b/B} \end{bmatrix} \begin{bmatrix} {}^{b}\dot{\boldsymbol{v}}_{B/e} \\ {}^{b}\dot{\boldsymbol{\omega}}_{b/e} \end{bmatrix} = \begin{bmatrix} \boldsymbol{f}_b - \boldsymbol{\omega}_{b/e} \times \boldsymbol{p}_{b/e} \\ \boldsymbol{g}_{b/B} - \boldsymbol{\omega}_{b/e} \times \boldsymbol{h}_{b/B} \end{bmatrix} \tag{2.18}$$

B.4.3 Model 6c

Another option is to target ${}^{b}\dot{\boldsymbol{v}}_{RM/e}$ directly, but again using the momentum about the body center of mass B. Like *Model 6b* this also produces a simpler dynamics model, but again at the cost of making it less convenient to precompute the apparent inertia matrix.

The main purpose of this model is for validating model implementations. An implementation of this model is available as `Paraglider6c` in the `glidersim` package.

Computing the inertial derivatives with respect to the body frame:

$$\begin{aligned} {}^{e}\dot{\boldsymbol{p}}_{b/e} &= m_b \left({}^{b}\dot{\boldsymbol{v}}_{RM/e} + {}^{b}\dot{\boldsymbol{\omega}}_{b/e} \times \boldsymbol{r}_{B/RM} \right) + \boldsymbol{\omega}_{b/e} \times \boldsymbol{p}_{b/e} \\ {}^{e}\dot{\boldsymbol{h}}_{b/B} &= \boldsymbol{J}_{b/B} \cdot {}^{b}\dot{\boldsymbol{\omega}}_{b/e} + \boldsymbol{\omega}_{b/e} \times \boldsymbol{h}_{b/B} \end{aligned} \tag{2.19}$$

Using the body center of mass as the reference point simplifies the equation for angular momentum:

$$\begin{aligned} {}^{e}\dot{\boldsymbol{p}}_{b/e} &= \boldsymbol{f}_b \\ {}^{e}\dot{\boldsymbol{h}}_{b/B} &= \boldsymbol{g}_{b/B} \end{aligned} \tag{2.20}$$

Combining (2.19) and (2.20): and rewriting as a linear system:

$$\begin{bmatrix} m_b & -m_b \left[\boldsymbol{r}_{B/RM} \right]^{\times} \\ 0 & \boldsymbol{J}_{b/B} \end{bmatrix} \begin{bmatrix} {}^{b}\dot{\boldsymbol{v}}_{RM/e} \\ {}^{b}\dot{\boldsymbol{\omega}}_{b/e} \end{bmatrix} = \begin{bmatrix} \boldsymbol{f}_b - \boldsymbol{\omega}_{b/e} \times \boldsymbol{p}_{b/e} \\ \boldsymbol{g}_{b/B} - \boldsymbol{\omega}_{b/e} \times \boldsymbol{h}_{b/B} \end{bmatrix} \tag{2.21}$$

B.4.4 Model 9a

Similar to *Model 6a*, this design uses the riser connection midpoint *RM* as the reference point for both the body and the payload, which simplifies incorporating the apparent mass matrix. However, this model treats the body and payload as separate components, connected by a rotational spring-damper model that adds an additional three degrees-of-freedom. A similar 9DoF model derivation can be found in [35] (9DoF, but relative roll and pitch are unconstrained).

An implementation of this model is available as `Paraglider9a` in the `glidersim` package. The `glidersim` package also includes `Paraglider9b`, which uses the centers of mass as the reference points for the body and payload dynamics; that choice simplifies the derivatives for angular momentum (because it eliminates the moment arms), but prohibits incorporating the effects of apparent mass.

B.4.4.1 Real mass only Start with the equations for the translational and angular momentum of the body b about the reference point RM as observed by the inertial reference frame e:

$$\begin{aligned} \boldsymbol{p}_{b/e} &= m_b \, \boldsymbol{v}_{B/e} \\ &= m_b \left(\boldsymbol{v}_{RM/e} + \boldsymbol{\omega}_{b/e} \times \boldsymbol{r}_{B/RM} \right) \end{aligned} \tag{2.22}$$

$$\begin{aligned} \boldsymbol{p}_{p/e} &= m_p \, \boldsymbol{v}_{P/e} \\ &= m_p \left(\boldsymbol{v}_{RM/e} + \boldsymbol{\omega}_{b/e} \times \boldsymbol{r}_{P/RM} \right) \end{aligned} \tag{2.23}$$

$$\boldsymbol{h}_{b/RM} = m_b \, \boldsymbol{r}_{B/RM} \times \boldsymbol{v}_{RM/e} + \boldsymbol{J}_{b/RM} \cdot \boldsymbol{\omega}_{b/e} \tag{2.24}$$

$$\boldsymbol{h}_{p/RM} = m_p \, \boldsymbol{r}_{P/RM} \times \boldsymbol{v}_{RM/e} + \boldsymbol{J}_{p/RM} \cdot \boldsymbol{\omega}_{p/e} \tag{2.25}$$

Compute the two momentum derivatives:

$$\begin{aligned} {}^{e}\dot{\boldsymbol{p}}_{b/e} &= {}^{b}\dot{\boldsymbol{p}}_{b/e} + \boldsymbol{\omega}_{b/e} \times \boldsymbol{p}_{b/e} \\ &= m_b \left({}^{b}\dot{\boldsymbol{v}}_{RM/e} + {}^{b}\dot{\boldsymbol{\omega}}_{b/e} \times \boldsymbol{r}_{B/RM} \right) + \boldsymbol{\omega}_{b/e} \times \boldsymbol{p}_{b/e} \\ {}^{e}\dot{\boldsymbol{h}}_{b/RM} &= {}^{b}\dot{\boldsymbol{h}}_{b/RM} + \boldsymbol{\omega}_{b/e} \times \boldsymbol{h}_{b/RM} \\ &= m_b \boldsymbol{r}_{B/RM} \times {}^{b}\dot{\boldsymbol{v}}_{RM/e} + \boldsymbol{J}_{b/RM} \cdot {}^{b}\dot{\boldsymbol{\omega}}_{b/e} + \boldsymbol{\omega}_{b/e} \times \boldsymbol{h}_{b/RM} \\ {}^{e}\dot{\boldsymbol{p}}_{p/e} &= {}^{p}\dot{\boldsymbol{p}}_{p/e} + \boldsymbol{\omega}_{p/e} \times \boldsymbol{p}_{p/e} \\ &= m_p \left({}^{p}\dot{\boldsymbol{v}}_{RM/e} + {}^{p}\dot{\boldsymbol{\omega}}_{p/e} \times \boldsymbol{r}_{P/RM} \right) + \boldsymbol{\omega}_{p/e} \times \boldsymbol{p}_{p/e} \\ &= m_p \left({}^{b}\dot{\boldsymbol{v}}_{RM/e} + \boldsymbol{\omega}_{b/p} \times \boldsymbol{v}_{RM/e} + {}^{p}\dot{\boldsymbol{\omega}}_{p/e} \times \boldsymbol{r}_{P/RM} \right) + \boldsymbol{\omega}_{p/e} \times \boldsymbol{p}_{p/e} \\ {}^{e}\dot{\boldsymbol{h}}_{p/RM} &= {}^{p}\dot{\boldsymbol{h}}_{p/RM} + \boldsymbol{\omega}_{p/e} \times \boldsymbol{h}_{p/RM} \\ &= m_p \boldsymbol{r}_{P/RM} \times {}^{p}\dot{\boldsymbol{v}}_{RM/e} + \boldsymbol{J}_{p/RM} \cdot {}^{p}\dot{\boldsymbol{\omega}}_{p/e} + \boldsymbol{\omega}_{p/e} \times \boldsymbol{h}_{p/RM} \\ &= m_p \boldsymbol{r}_{P/RM} \times \left({}^{b}\dot{\boldsymbol{v}}_{RM/e} + \boldsymbol{\omega}_{b/p} \times \boldsymbol{v}_{RM/e} \right) + \boldsymbol{J}_{p/RM} \cdot {}^{p}\dot{\boldsymbol{\omega}}_{p/e} + \boldsymbol{\omega}_{p/e} \times \boldsymbol{h}_{p/RM} \end{aligned} \tag{2.26}$$

Derivatives of the payload momentums are computed in terms of the body velocity derivative in the body frame to allow writing the dynamics as a single system of equations. First,

compute the net external forces and moments:

$$\begin{aligned} \boldsymbol{f}_b &= \boldsymbol{f}_{b,\text{aero}} + \boldsymbol{f}_{b,\text{weight}} \\ \boldsymbol{g}_{b/RM} &= \boldsymbol{g}_{b,\text{aero}} + \boldsymbol{g}_{b,\text{weight}} \\ \boldsymbol{f}_p &= \boldsymbol{f}_{p,\text{aero}} + \boldsymbol{f}_{p,\text{weight}} \\ \boldsymbol{g}_{p/RM} &= \boldsymbol{g}_{p,\text{aero}} + \boldsymbol{g}_{p,\text{weight}} \end{aligned} \tag{2.27}$$

And equate them to the derivatives of momentum with respect to the inertial frame:

$$\begin{aligned} {}^e\dot{\boldsymbol{p}}_{b/e} &= \boldsymbol{f}_b - \boldsymbol{f}_{RM} \\ {}^e\dot{\boldsymbol{h}}_{b/RM} + \boldsymbol{v}_{RM/e} \times \boldsymbol{p}_{b/e} &= \boldsymbol{g}_{b/RM} - \boldsymbol{g}_{RM} \\ {}^e\dot{\boldsymbol{p}}_{p/e} &= \boldsymbol{f}_p + \boldsymbol{f}_{RM} \\ {}^e\dot{\boldsymbol{h}}_{p/RM} + \boldsymbol{v}_{RM/e} \times \boldsymbol{p}_{p/e} &= \boldsymbol{g}_{p/RM} + \boldsymbol{g}_{RM} \end{aligned} \tag{2.28}$$

The spring-damper connection produces forces and moments shared by the body and the payload. There are six variables but only three degrees of freedom. Both systems have the riser connection point RM at a fixed position, and the force only exists to maintain the fixed relative positioning.

$$\boldsymbol{g}_{RM} = \begin{bmatrix} \kappa_\phi \phi + \kappa_{\dot{\phi}} \dot{\phi} \\ \kappa_\theta \theta + \kappa_{\dot{\theta}} \dot{\theta} \\ \kappa_\gamma \gamma + \kappa_{\dot{\gamma}} \dot{\gamma} \end{bmatrix} \tag{2.29}$$

Where $\boldsymbol{\omega}^p_{p/b} = \langle \phi, \theta, \gamma \rangle$ are the angular rates of the payload, ${}^p\dot{\boldsymbol{\omega}}^p_{p/b} = \left\langle \dot{\phi}, \dot{\theta}, \dot{\gamma} \right\rangle$ are the angular accelerations of the payload, and the κ are the stiffness and dampening coefficients of the spring-damper model.

This is a very simple model. A better model would need to account for the coupling between dimensions, and should really be a function of the riser strap width.]]

Combining equations (2.26) and (2.28) and rewriting as a linear system provides the dynamics of the real mass (solid mass plus the enclosed air) in terms of ${}^b\dot{\boldsymbol{v}}_{RM/e}$, ${}^b\dot{\boldsymbol{\omega}}_{b/e}$,

${}^b\dot{\boldsymbol{\omega}}^p_{p/e}$, and $\boldsymbol{f}^b_{RM}$:

$$\boldsymbol{A}_{r/RM}\begin{bmatrix} {}^b\dot{\boldsymbol{v}}^b_{RM/e} \\ {}^b\dot{\boldsymbol{\omega}}^b_{b/e} \\ {}^p\dot{\boldsymbol{\omega}}^p_{p/e} \\ \boldsymbol{f}^b_{RM} \end{bmatrix} = \begin{bmatrix} \boldsymbol{b}^b_1 \\ \boldsymbol{b}^b_2 \\ \boldsymbol{b}^p_3 \\ \boldsymbol{b}^p_4 \end{bmatrix} \tag{2.30}$$

Where:

$$\boldsymbol{A}_{r/RM} = \begin{bmatrix} m_b\,\boldsymbol{I}_3 & -m_b\left[\boldsymbol{r}^b_{B/RM}\right]^\times & \boldsymbol{0}_{3\times3} & \boldsymbol{I}_3 \\ m_b\left[\boldsymbol{r}^b_{B/RM}\right]^\times & \boldsymbol{J}^b_{b/RM} & \boldsymbol{0}_{3\times3} & \boldsymbol{0}_{3\times3} \\ m_p\,\boldsymbol{C}_{p/b} & \boldsymbol{0}_{3\times3} & -m_p\left[\boldsymbol{r}^p_{P/RM}\right]^\times & -\boldsymbol{C}_{p/b} \\ m_p\left[\boldsymbol{r}^p_{P/RM}\right]^\times\boldsymbol{C}_{p/b} & \boldsymbol{0}_{3\times3} & \boldsymbol{J}^p_{p/RM} & \boldsymbol{0}_{3\times3} \end{bmatrix}$$

$$\begin{aligned}
\boldsymbol{b}^b_1 &= \boldsymbol{f}^b_b - \boldsymbol{\omega}^b_{b/e}\times\boldsymbol{p}^b_{b/e} \\
\boldsymbol{b}^b_2 &= \boldsymbol{g}^b_b - \boldsymbol{g}^b_{RM} - \boldsymbol{v}^b_{RM/e}\times\boldsymbol{p}^b_{b/e} - \boldsymbol{\omega}^b_{b/e}\times\boldsymbol{h}^b_{b/RM} \\
\boldsymbol{b}^p_3 &= \boldsymbol{f}^p_p - \boldsymbol{\omega}^p_{p/e}\times\boldsymbol{p}^p_{p/e} - m_p\boldsymbol{\omega}^p_{b/p}\times\boldsymbol{v}^p_{RM/e} \\
\boldsymbol{b}^p_4 &= \boldsymbol{g}^p_b + \boldsymbol{g}^p_{RM} - \boldsymbol{v}^p_{RM/e}\times\boldsymbol{p}^p_{p/e} - \boldsymbol{\omega}^p_{p/e}\times\boldsymbol{h}^p_{p/RM} - m_p\boldsymbol{r}^p_{P/RM}\times\left(\boldsymbol{\omega}^p_{b/p}\times\boldsymbol{v}^p_{RM/e}\right)
\end{aligned} \tag{2.31}$$

B.4.4.2 Real mass + apparent mass As with the 6-DoF system, the effects of apparent mass on the canopy can be accounted for by adding the apparent inertia matrix from *Apparent mass of a parafoil* to the components of the system matrix associated with the translational and angular acceleration of the body and accounting for the translational and angular apparent momentum:

$$\left(\boldsymbol{A}_{r/RM} + \begin{bmatrix} \boldsymbol{A}_{a/RM} & \boldsymbol{0}_{6\times6} \\ \boldsymbol{0}_{6\times6} & \boldsymbol{0}_{6\times6} \end{bmatrix}\right)\begin{bmatrix} {}^b\dot{\boldsymbol{v}}^b_{RM/e} \\ {}^b\dot{\boldsymbol{\omega}}^b_{b/e} \\ {}^p\dot{\boldsymbol{\omega}}^p_{p/e} \\ \boldsymbol{f}^b_{RM} \end{bmatrix} = \begin{bmatrix} \boldsymbol{b}^b_5 \\ \boldsymbol{b}^b_6 \\ \boldsymbol{b}^p_3 \\ \boldsymbol{b}^p_4 \end{bmatrix} \tag{2.32}$$

$$\begin{aligned}
\boldsymbol{b}^b_5 &= \boldsymbol{b}^b_1 - \boldsymbol{\omega}_{b/e}\times\boldsymbol{p}_{a/e} \\
\boldsymbol{b}^b_6 &= \boldsymbol{b}^b_2 - \boldsymbol{v}_{RM/e}\times\boldsymbol{p}_{a/e} - \boldsymbol{\omega}_{b/e}\times\boldsymbol{h}_{a/RM} + \boldsymbol{v}_{RM/e}\times\left(\boldsymbol{M}_a\cdot\boldsymbol{v}_{RM/e}\right)
\end{aligned}$$

Where $\boldsymbol{A}_{a/RM}$ is the apparent inertia matrix of the canopy from (2.4), $\boldsymbol{M}_a$ is the apparent mass matrix from (2.2), and $\boldsymbol{p}_{a/e}$ and $\boldsymbol{h}_{a/RM}$ are the linear and angular apparent momentums from (2.5) and (2.6). The extra term $\boldsymbol{v}_{RM/e} \times \left(\boldsymbol{M}_a \boldsymbol{v}_{RM/e}\right)$ in $\boldsymbol{b}_6^b$ is necessary to avoid double counting the aerodynamic moment already accounted for by the section pitching coefficients.